



**US Army Corps
of Engineers**
Hydrologic Engineering Center

Developing Seasonal and Long-term Reservoir System Operation Plans using HEC-PRM

June 1996

19961018 073

DTIC QUALITY INSPECTED 3

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204 Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED	
	June 1996	Research Document 40, June 1996	
4. TITLE AND SUBTITLE		5. FUNDING NUMBERS	
Developing Seasonal and Long-term Reservoir System Operation Plans using HEC-PRM		WU 32939	
6. AUTHOR(S)			
Jay R. Lund, under contract to Hydrologic Engineering Center			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)		8. PERFORMING ORGANIZATION REPORT NUMBER	
Hydrologic Engineering Center U.S. Army Corps of Engineers 609 Second Street Davis, CA 95616-4687		RD-40	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
N/A		N/A	
11. SUPPLEMENTARY NOTES			
N/A			
12a. DISTRIBUTION / AVAILABILITY STATEMENT		12b. DISTRIBUTION CODE	
Distribution of this document is unlimited.			
13. ABSTRACT (Maximum 200 words)			
<p>The major intent of this report is to provide guidance for the use of HEC-PRM (and deterministic optimization models in general) for seasonal and long-term operation planning. Seasonal and long-term operation planning are fairly different undertakings, in the level of analysis detail, the role of hydrology and uncertainty, and the frequency of policy revision (seasonal operations can be revised monthly). Thus the details of setting up, conducting, and interpreting optimization model results can differ significantly between these two applications. For both optimization model applications, simulation modeling is an important, if not essential, companion. As such, it is useful to review the wide variety of operating-rule forms available for reservoir system simulation modeling.</p>			
14. SUBJECT TERMS		15. NUMBER OF PAGES	
HEC Prescriptive Reservoir Model (HEC-PRM) optimization, deterministic, seasonal operation rules, long-term operational rules, penalty functions, streamflow, forecasts, floods, droughts, environmental, hydropower		125	
16. PRICE CODE			
17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	UL



REPLY TO
ATTENTION OF

CEWRC-HEC-T

DEPARTMENT OF THE ARMY
CORPS OF ENGINEERS, WATER RESOURCES SUPPORT CENTER
HYDROLOGIC ENGINEERING CENTER
609 SECOND STREET
DAVIS, CALIFORNIA 95616-4687

11 October 1996

MEMORANDUM FOR DIR USADLA, DTIC, ATTN: FDAC, 8725 John J. Kingman Road,
Suite 0944, Fort Belvoir, VA 22060-6218

SUBJECT: Submission of Publications

1. Enclosed is a DTIC card and hardcopy for each publication listed below for submission into the Technical Reference File:
 - a. CPD-5Q, HEC-5Q Appendix on WQA, Sep 86, 347 pp.
 - b. CPD-6, HEC-6, User's Manual, Aug 93, 286 pp.
 - c. CPD-66, HEC-UNET User's Manuals Jul 96, 288 pp.
 - d. CPD-67, HEC-PBA User's Manual, Sep 94, 154 pp.
 - e. PR-21, Columbia River System Analysis Model: Phase II, Dec 93, 176 pp.
 - f. PR-26, Preliminary Operating Rules for the Columbia River System from HEC-PRM Results, Jun 95, 147 pp.
 - g. RD-39, Water Control Data System(WCDS): Past, Present and Future, Sep 95, 100 pp.
 - h. RD-40, Developing Seasonal and Long-term Reservoir System Operation Plans using HEC-PRM, Jun 96, 125 pp.
 - i. RD-43, Application of HEC-PRM for Seasonal Reservoir Operation of the Columbia River System, Jun 96, 188 pp.
 - j. RD-44, Closures and Interior Facilities for Levee Projects; Principles, Case Examples, and Risk-based Analysis Concepts, Sep 96, 96 pp.
 - k. TP-155, Runoff Simulation Using Radar Rainfall Data, Aug 96, 8 pp.
2. Please provide reference numbers ASAP. For questions, POC is Ms. Josie Garcia-Moreno.


For DARRYL W. DAVIS
Director

Encls

Developing Seasonal and Long-term Reservoir System Operation Plans using HEC-PRM

June 1996

Hydrologic Engineering Center
U.S. Army Corps of Engineers
609 Second Street
Davis, CA 95616-4687

(916) 756-1104

RD-40

Table of Contents

Chapter	Page
List of Figures	iii
List of Tables	iv
Preface	v
Summary	vi
1 Introduction	1
1.1 Problem and Approach	1
1.2 Reservoir Operation Studies	1
1.3 Simulation and Optimization	2
1.4 Evaluation of Reservoir Operating Rules	2
1.5 Operating Objectives and Penalty Functions	3
1.6 HEC-PRM	4
1.7 Implicit Stochastic Optimization	5
1.8 Overview of Report	6
2 Seasonal and Long-term Reservoir Operating	7
2.1 Introduction	7
2.2 Single Reservoir Rules	8
2.3 Rules for Reservoirs in Series	31
2.4 Rules for Reservoirs in Parallel	37
2.5 Miscellaneous Multi-Reservoir Rules	44
2.6 Decomposition of Multi-Reservoir Systems	49
2.7 Conclusions	49
Appendix	
2.1: Derivation of Hydropower Production Rules for Reservoirs in Series ..	51
2.2: Derivation of Space Rules	53
2.3: "Storage Effectiveness Index" Rules for Hydropower Production	57
3 Long-term Rules from Optimization	59
3.1 Introduction	59
3.2 Hydrologies for Rule Development	59
3.3 Inferring Rules from Results	62
3.4 Mixed Simulation-Optimization Approaches	68
3.5 Testing and Refining Operating Rules by Simulation	70
3.6 Some Pitfalls in Operating Rule Development	71
3.7 Conclusions	73
4 Seasonal Rules from Optimization Studies	75
4.1 Introduction	75
4.2 Penalty Functions and Initial and End-of-period Conditions	75
4.3 Hydrologic Input	78

Table of Contents (continued)

Chapter		Page
4.4	Operational Uses of Many-Event Runs	80
4.5	Development of Annual Operating Rule Curves	84
4.6	Testing and Refining Operating Rules by Simulation	85
4.7	Seasonal Operation Update Studies	86
4.8	Management Uses of Many-Event Results	86
4.9	Conclusions	89
5	Conclusions	91
Appendices		
A	References	A-1
B	Summary of Columbia River System Seasonal Operation Study	B-1

List of Figures

Figure	Page
2.1 Standard Operating Policy	10
2.2 Hedging Release Rule	11
2.3 P1-P2 Hedging Rules	13
2.4 Trigger Value (One Point) Hedging	14
2.5 Zone Based Hedging	16
2.6 Optimal Hedging Policy from Stochastic Dynamic Program	17
2.7 Value of Release	18
2.8 Flood Control Storage Fill and Evacuation Rules	20
2.9 Missouri River System Zone Based Operation	23
2.10 Zone Based Operation	25
2.11 Arkansas Zone-based Operation	26
2.12 Storage Target Rule and Actual Operations	28
2.13 Inflow Locations Involved in an Oahe Reservoir Release Rule	30
2.14 Reservoirs in Series	31
2.15 Change in Head With Varying Capacities	35
2.16 Reservoirs in Parallel	37
2.17 Storage Allocation Plot for Three Reservoirs in the Missouri River System for July	46
2.18 Decomposition of a Complex Multi-Reservoir System into a Sub-System in Series and a Sub-System in Parallel	49
3.1 Quartiles of Seasonal Storage in Fort Randall under HEC-PRM Operations	64
3.2 July Storage Allocation Results for the Missouri River System	64
3.3 July Relationship between Systemwide Storage and System Release for the Missouri River System	66
3.4 Comparison of Systemwide Storage Between HEC-PRM and HYSSR for the Columbia River System	67
4.1 End-of-period Storage Penalties	77
4.2 Example Storage Position Analysis Plot	80
4.3 Quartile Plot of Storages from Position Analysis Plot in Figure 4.2	81
4.4 Example Position Analysis Results for Mica Reservoir, Columbia R. System, for January-July, 1994	81
4.5 Example Position Analysis Results for Arrow Reservoir, Columbia R. System, for January-July, 1994	82
4.6 Example Storage Quartile Results for Arrow Reservoir, Columbia R. System, for January-July, 1994	82
4.7 Storage Allocation Plot for Variable Drawdown Season, January-March, 1995 for Arrow, Duncan, and Libby Reservoirs in the Columbia River System	83
4.8 Monthly Exceedence Probability Plots for Storage for March through July	87
4.9 Example Non-Exceedence Probability Plots for Water Yield, May-July	88
4.10 Example Non-Exceedence Probability Plots for Monthly Power Yields for June-December	88

List of Tables

Table		Page
1.1	Modeling Characteristics for Reservoir Operations Problems	2
2.1	Typically Desirable Single-Purpose and Common Multi-Purpose Reservoir Operating Rules (by section number)	7
2.2	Hedging in Occoquan Reservoir Operations	13
2.3	"Index level" Operation Table for Storage Allocation Plot in Figure 2.15	48
3.1	Rationale and Uses for Simulation Modeling in Optimization Rule-Making	70
3.2	Summary of Approaches for Developing Long-Term Operating Policies from Optimization Model Results	73
4.1	Summary of Approaches for Developing Annual Operating Rules from Optimization Model Results	84
4.2	Indicators of Performance For Various Reservoir System Purposes	87

Preface

The central problem discussed in this report is the operation of reservoir systems. This is a classical problem. Reservoir operations theory had its formal beginnings over seventy-five years ago (Varlet, 1923). Computerized simulation modeling has been applied to water resource planning and operation studies for over forty years (a Corps of Engineers Missouri River system study in 1953, cited in Hall and Dracup, 1970). Numerical optimization models have been studied for application to water resource systems for over thirty years (Buras, 1963). While a great deal has been learned about how to operate complex reservoir systems for a wide variety of purposes and for multiple purposes, we have not solved this problem. Indeed, there are reasons to think that this problem cannot be completely solved, given the uncertainties involved, the common presence of multiple near-optimal solutions, and the conflicts inherent in making trade-offs between multiple system purposes (Miquel and Roche, 1983; Rogers and Fiering, 1986).

Despite our lack of perfect solutions, reservoir operations remains an important engineering problem in the economic, environmental, and political arena. Recent changes in the emphasis placed on different reservoir system purposes locally and nationally have caused the profession to undertake unprecedented large-scale studies focusing largely on the operation (as opposed to the capacity planning) of existing reservoirs. These have occurred for most of the large systems operated by the Corps of Engineers, such as the Columbia River system, the Missouri River System, and the Apalachicola-Chattahoochee-Flint-Alabama-Coosa-Tallapoosa (ACT-ACF) river systems.

The traditional method in practice for conducting reservoir operations studies is through simulation modeling, using a simulation model to examine and evaluate a variety of operating policies. The recent Missouri River study examined over 104 operating policies. However, even such large numbers of potential policies cannot reflect the wide and multi-dimensional variety of operating rules available to planners.

A second practical approach for finding desirable reservoir operation policies is available as a companion to purely simulation based approaches. This is the application of deterministic optimization models (typically network flow models such as HEC-PRM) to find promising sets of rules as a point of departure for more detailed simulation studies. This report provides theoretical and practical background for the application of this approach to real reservoir systems. The guidance in this report results in part from the lessons learned from applying HEC-PRM to the Columbia River System (USACE 1993a, 1995, 1996), the Missouri River mainstem system (USACE, 1992, 1994b), and Alamo Reservoir in Arizona (Kirby, 1994). A wide variety of results from academic and other practical studies are also useful in the application and interpretation of deterministic optimization models to reservoir operation problems. This report focuses on the application of deterministic optimization models to long-term/strategic operating rule studies and annual/seasonal operating rule studies. To further these applications a fairly detailed review of potential reservoir operating rules available for interpreting deterministic optimization results is provided.

This report was written principally by Professor Jay R. Lund of the Department of Civil and Environmental Engineering at the University of California, Davis. Joel Guzman, an undergraduate student in Civil and Environmental Engineering, completed the initial draft of Chapter 2. This work was supervised by Mike Burnham and Darryl Davis, and was aided considerably by suggestions from Dan Barcellos and Richard Hayes, U.S. Army Corps of Engineers Hydrologic Engineering Center. Morris Israel, Ken Kirby, and Mimi Jenkins also provided valuable comments on this material.

Summary

This report is motivated by two problems, one new and one traditional. First, a continuing traditional problem for reservoir operations is seasonal or annual operations planning, attempting to tailor long-term operation policies for current conditions of storage, streamflow forecasts, and facility conditions (such as short-term turbine maintenance schedules). Many reservoir systems face such problems every year; some face these problems every month or every day. During drought or flood conditions, most system operators seek improved operations by considering appropriate improvements over long-term operating policies.

The second problem is the changing purposes, regulation conditions, and public interests for the management of reservoir systems. Most of these changes have originated with increased environmental concerns, but significant changes also have occurred in recreational, water supply, and other more traditional operating purposes. These changes in societal objectives for reservoir management have prompted re-examination and re-regulation of many (if not most) reservoir systems across the country. These re-examinations often have been accompanied by heightened levels of controversy and technical scrutiny. This report provides guidance for the use of relatively new reservoir management software to address these problems.

The Hydrologic Engineering Center's Prescriptive Reservoir Model (HEC-PRM) is a prescriptive, optimization, or operating purpose-driven model for reservoir system operations. The model suggests promising reservoir operation decisions, driven by quantitative descriptions of operating purpose values, called penalty functions. This represents a new, but well-tested technology, which (together with other optimization models) has become an increasingly popular tool for reservoir system studies.

The major intent of this report is to provide guidance for the use of HEC-PRM (and deterministic optimization models in general) for seasonal and long-term operation planning. Seasonal and long-term operation planning are fairly different undertakings, in the level of analysis detail, the role of hydrology and uncertainty, and the frequency of policy revision (seasonal operations can be revised monthly). Thus the details of setting up, conducting, and interpreting optimization model results can differ significantly between these two applications. For both optimization model applications, simulation modeling is an important, if not essential, companion. As such, it is useful to review the wide variety of operating rule forms available for reservoir system simulation modeling.

Use of HEC-PRM for Seasonal Operations

Where there is value in modifying reservoir operating policies to account for specific current conditions, such as current reservoir levels, flow forecasts, and short-term variation in system operating purpose values, annual or seasonal operation plans are commonly developed. These operation plans typically are based on long-term operation policies, often modified or adapted based on system simulation studies. These simulation modeling studies are frequently rather intensive affairs, requiring many simulation runs to examine a range of alternative operating plans for the coming year or season.

HEC-PRM, as an optimization model, can quickly identify promising operation strategies for seasonal operations. To do this, HEC-PRM usually must make several simplifications of the system compared with most simulation models. Therefore, simulation studies remain necessary to refine and test the operation plans suggested by HEC-PRM results. However, the use of HEC-PRM should allow these simulation studies to be more focused. Moreover, since HEC-PRM is driven to maximize a system's operating purpose values, the resulting operation plan will have more rigorous economic support.

Several special technical problems must be addressed in setting up a HEC-PRM model for an annual or seasonal operating plan application. These include establishing initial storage conditions, setting penalties and targets for end-of-period storages, modifying penalty functions to reflect short-term conditions (such as turbine maintenance schedules), and representation of likely future inflows. A variety of forms of annual and seasonal operating advice can be found in HEC-PRM results. These include suggestions for refill probabilities, short-term release and storage advice, seasonal storage targets, drawdown and refill ordering (for multi-reservoir systems), and yield-reliability results. These are discussed here and have been applied preliminarily to the Columbia River System (USACE, 1996).

Developing Long-Term Operating Rules using HEC-PRM

The second problem discussed above, changing long-term reservoir operation policies to reflect changes in reservoir operating purposes, often requires a rethinking of reservoir operation strategies. Simulation modeling retains a central role in this process, but faces difficulties due to the enormous range of different operating policies which must be examined in these cases. Again, HEC-PRM can be used to suggest the most promising reservoir operation strategies, which then can be refined and tested using simulation studies (USACE, 1992, 1994b). HEC-PRM can also be used as a screening model to examine changes in facility capacities (USACE, 1993a), changes in operating purposes (USACE, 1993a), and potential conflicts among reservoir users (Kirby, 1994).

A variety of approaches are available for developing operating rules from HEC-PRM results. These are summarized in Table 1. Frequently, several of these approaches will be used together, with some approaches identifying the most important variables, others suggesting a promising structure for an operating policy, and still other approaches employed for refining and testing an operating policy (USACE, 1992).

Table 1: Summary of Approaches for Developing Operating Policies from Optimization Model Results

Rule Development Approach	Summary Comments
Direct use of results	Unlikely to be possible for long-term operation planning studies.
Intuitive Approaches	Probably of some use in all cases.
Operation rule theory	Useful for developing operation strategy, but probably insufficient alone.
Regression	Often difficult for real applications, but often can identify important variables.
Principal components analysis	May work well in some cases.
Artificial neural networks (ANN)	Requires unusually large amounts of streamflow data.
Incremental modification of existing rules	This should be the most promising approach, given good existing rules, small changes in operating purposes, and an existing simulation model.
Mixed simulation-optimization	Probably essential in most cases.

The Importance of Simulation Modeling

HEC-PRM and other optimization models are unlikely to be able to replace simulation models for the operation of major reservoir systems within the near future. Optimization models still require important simplifications of the system in order to be solvable, such as smoothing some peaking operations, simplifying the system configuration in some ways, and accepting only some forms of penalty functions for operating purposes. Simulation models remain necessary to examine the detailed performance of system operation, although simulation studies can benefit from the focus and external validation provided by optimization model results.

Review of Operating Rule Forms

Since simulation will remain an essential part of reservoir system operation and a great deal of the use of optimization model results is to suggest operating rules for simulation studies, it is useful to review the wide range of operating rule forms available for reservoir operations. Over the last century or more, a variety of reservoir operating rules have been developed for particular reservoir purposes and reservoir configurations. These rules provide a great deal of flexibility in the specification of system operations under various flow, storage, and demand conditions.

A particular set of operating rules can be supported technically in a number of ways. Some rules are based on simple engineering principles for reservoir operations, such as keeping reservoirs full for water supply or empty for flood control. Several rule forms can be derived from formal optimization principles, including several rules for allocating stored water among reservoirs, such as the New York City space rules and hydropower production and energy storage rules. However, many rules are based largely on empirical or experimental successes, either from actual operational performance, performance in simulation studies, or optimization results. These experimentally-supported rules are likely to be the most common for multi-purpose projects.

Many opportunities exist for the use of formal optimization methods within reservoir simulation models. Examples include implementing storage allocation rules for reservoirs in series and in parallel, as well as general penalty-minimizing operations for allocating water among storage, flow, and diversion uses within a given time-step. In these applications, the intent is to model an optimizing system operator who seeks to maximize system performance.

Chapter 1

Introduction

1.1 Problem and Approach

The operation of reservoirs is a classical problem in water resources engineering with a long and distinguished history, spanning well over a century (Rippl, 1883; Varlet, 1923; Maass et al, 1966; USACE, 1991d). While much recent work on this problem has become very sophisticated, especially in dealing with multiple-purpose reservoirs, the problem remains unresolved technically. We do not know exactly how to operate multiple purpose reservoirs when faced with hydrologic and other uncertainties.

This report examines one practical approach to this problem, combining the use of prescriptive (optimization) modeling with more common simulation modeling techniques. While the results of this approach are not guaranteed to be an "optimal" solution for reservoir operation, the results should offer very good solutions with a great deal of formal economic justification (NED-benefit maximization) and a relatively transparent methodology, at relatively little overall study expense.

The gist of the approach is to employ a deterministic optimization model (such as HEC-PRM) to provide one or more promising sets of preliminary rules for reservoir operation. These rules are then tested and refined with much more detailed, realistic, and common reservoir system simulation models. This report examines how this approach can and has been applied to both seasonal and long-term operation problems.

1.2 Reservoir Operation Studies

The overall management of reservoir and water resource systems is summarized by USACE (1987). Within this larger and more detailed context of reservoir management, reservoir operations studies are used to identify promising operational policies for various time-scales. The three most common settings for reservoir operation studies are:

- long-term strategic reservoir operations,
- seasonal operations studies, and
- real time operations.

Long term or strategic operations studies focus on general guidance and targets for system operations, including seasonal storage targets, establishment of general policies for flood storage volumes and their general seasonal operations, and drought operations. Such studies are typically made using the historical record, the drought of record, and major historical or synthetic floods to establish general long-term operating policies. Such guidance typically is found in the Master Water Control Manuals for Corps reservoirs. Long-term reservoir operation studies are conducted rather infrequently and may require a year or more to complete.

Seasonal operation studies address more particular and nearer-term conditions within the context of long-term operations. Examples of seasonal operations studies include the annual operation studies conducted for many systems and the system refill and drawdown studies conducted on a seasonal basis for some systems. Seasonal operation studies are conducted for every operating season (annual, semi-annual, or monthly) for a particular reservoir system. These studies must be completed relatively quickly, often within a period ranging from a week to a month, and so must become rather routine.

Real time operation studies address very short term operations within the context of seasonal operations plans. Real time studies involve both routine short-term hydropower and water delivery scheduling and emergency operations for flood control or drought. Real time operation planning for hydropower and water supply scheduling is a daily or weekly exercise for many systems. For emergency operations such as flood control or drought, studies must be conducted very quickly, often within a few hours or less for flood control problems, and are not routine, but must be done with great forethought and experience. Real time flood, hydropower, and environmental operations typically must take place on an hourly or daily time-step, often with special attention to the smooth transition of releases to reduce potential safety and habitat problems downstream, around the reservoir rim, and with dam and outlet structures. Also, river and streamflow forecasting are often much more accurate and valuable for real time operations.

Currently, most large and moderate-sized reservoir systems employ simulation modeling for all three types of studies. Some common model characteristics for these problems are summarized in Table 1.1. Real time hydropower, water delivery, and flood control modeling studies typically involve measuring performance at smaller time scales and often require representation of water routing lags, therefore requiring use of hourly or daily time-steps.

Table 1.1: Modeling Characteristics for Reservoir Operations Problems

Problem-Type	Frequency of Use	Period of Study Completion	Time-step
Long-Term	every 5-20 years	1-3 years	monthly, biweekly
Seasonal	several times per season	1-4 weeks	monthly, biweekly
Real Time	daily or episodically	1 hour to 7 days	daily, hourly

1.3 Simulation and Optimization

From its beginnings in the 1950s, simulation modeling has become the most common approach to finding and justifying reservoir operation policies for long-term, seasonal, and real time operating problems. The Corps led in the introduction of simulation modeling to reservoir problems, including the first Missouri River system model in 1953 (Hall and Dracup, 1970) and the introduction of early general-purpose reservoir simulation models in the 1960s (HEC-3 and HEC-5). Currently, most large reservoir systems use one or more simulation models for addressing long-term, seasonal, and real time operating problems. For example, several simulation models are used in the planning and management of different aspects of the Columbia River, Missouri River, California Central Valley, and Colorado River systems.

In most cases, simulation models require the explicit statement of operating rules. Operating rules, when integrated into a complete operating policy, are the instructions provided to the simulation model on how much water to release from each reservoir in a system under a complete range of possible conditions. Operating rules can take many forms, from IF-THEN statements to fairly complex mathematical functions, but, when combined into an operating policy, must always provide unambiguous operational instructions to the system simulation model. A wide range of reservoir operating rules is presented in Chapter 2, along with their theoretical underpinnings. In all cases, operating rules and policies must determine operating decisions (typically reservoir release decisions) based solely on conditions known to system operators, such as current month, current storage, current inflows, and sometimes available inflow forecasts.

Optimization models do not require explicit statements of operating rules, since operations are suggested (or prescribed) by the model. Instead, the objectives for reservoir operations must be explicitly stated in the form of *penalty functions*. In addition to the specification of quantitative

operating objectives (penalty functions), optimization models of reservoir systems also require *mathematical constraints* to represent physical, engineering, or legal constraints to the system and a representation of hydrologic inputs to the system. Physical and engineering constraints on the system would include reservoir capacities and turbine or outlet capacities. In some cases, minimum instream flows or other operational or legal constraints might be added, although it is usually preferable to represent such "soft" constraints with steep penalty functions. An optimization model, using often complex numerical solution algorithms, then prescribes desirable operating decisions which yield the minimum total penalty and satisfy all constraints defined for the system.

The representation of hydrologic uncertainty in optimization models has been developed from several different perspectives. The two important "schools of thought" on this subject are *explicitly stochastic* representation and *implicitly stochastic* representation. Explicitly stochastic representation of hydrologic uncertainty and variability requires characterizing the hydrologic inputs to the system in explicit probabilistic terms, i.e., joint probability functions and time-series correlations. These probability functions, disaggregated where more than one hydrologic input is important, are then typically formulated as a stochastic dynamic program and solved to derive an optimal operating policy (Stedinger, et al., 1984; Tejada-Guibert, et al., 1990, 1993, 1995). While rigorous, this explicit approach has significant computational inconvenience and requires often difficult specification of hydrologic input probability functions (Young, 1966). The alternative approach of implicitly stochastic optimization has lesser computational requirements, since it requires use of typically more efficient deterministic solution methods. Since HEC-PRM is a deterministic optimization model, this implicitly stochastic approach is the subject of the reservoir operating rule inference procedures presented in this report, and the subject of a more detailed section in this chapter.

In truth, it is likely that neither approach can identify the true optimal rules for reservoir operations. Unavoidable errors in formulation of the objective function, system constraints, and estimation of inflows, particularly under extreme conditions, make models imperfect representations of real systems (Miquel and Roche, 1983). However, operations found by either approach are likely to be "good," near-optimal, or promising given the wide range of near-optimal solutions that exist for many water resources problems (Rogers and Fiering, 1986).

1.4 Evaluation of Reservoir Operating Rules

Typically, the evaluation of operating rules is based on measures of performance included as part of the simulation model or as a separate post-processor (Hufschmidt and Fiering, 1966). Performance measures can be technical, such as firm yield (Rippl, 1883), reliability, resiliency, and vulnerability (Hashimoto, et al., 1982), or economic (Lund, et al., 1995). Comparison of reservoir operating rules given uncertainty in performance has been approached using stochastic dominance (Su, et al., 1991).

1.5 Operating Objectives and Penalty Functions

In the operation of a reservoir system, there is often competition between reservoir purposes, and sometimes competition within a reservoir purpose such as where different locations compete for limited flood control or water supply capability. Thus, it is important to develop methods of measuring how well a proposed set of reservoir operating policies is likely to perform on these often-competing objectives (USACE, 1977). The quantitative assessment of operating performance is most commonly done through the use of objective or penalty functions.

Objective function values usually are assessed directly from flow and storage results of the simulation model (Hufschmidt and Fiering, 1966; USACE, 1977). Thus, hydropower benefits

often are calculated based on hydropower production for each time-step and summed to provide an aggregate measure of hydropower benefits. It is most common for simple, static, objective functions to be used for assessing performance of operating purposes and objectives. These objectives typically take the form:

$$B_{jt} = f_{jt}(\vec{Q}_t, \vec{S}_t),$$

where

B_{jt} is the benefit to objective j at time t ,

$f_{jt}()$ is the benefit function for objective j at time t , it may vary daily or seasonally,

\vec{Q}_t is the vector of flows in the system at time t , and

\vec{S}_t is the vector of storages in the system at time t .

Overall performance for objective j is often estimated as the sum of performances for each period,

$$B_j = \sum B_{jt}$$

In other cases, the assessment of benefits or costs is more difficult, reflecting a need to examine longer-term dynamics in water management purpose demands and activities. Commonly, assessment of flood control benefits requires a separate dynamic model which can assess the dynamics of recovery from a series of flooding events resulting from a model. In this case, two floods occurring very close together often do much less damage than two floods spaced several years apart, when there is time for re-investment and recovery (USACE, 1994c). Similar long-term dynamics can be present for ecological and water supply performance (Lund, 1995). Ideally, each reservoir objective would be represented by a distinct objective or penalty function to allow explicit trade-offs of reservoir operating purposes (Kirby, 1994).

In practice, in-depth preparation of performance measures is not always possible and more expedient technical measures often are employed where more fundamental economic, social, or environmental performance measures are not readily available or involve potentially large estimation errors. Some common objective functions used to assess performance of reservoir systems include:

1. Minimize net economic losses/maximize net economic benefit,
2. Minimize weighted sum of objective penalties,
3. Maximize total usable water yield,
4. Maximize total energy yield,
5. Minimize peak flow,
6. Minimize mean annual water supply shortfall,
7. Minimize maximum flood peak, and
8. Maximize water supply reliability.

1.6 HEC-PRM

The Hydrologic Engineering Center Prescriptive Reservoir Model (HEC-PRM) is an optimization model which seeks to minimize a complex linear objective function for multi-reservoir systems subject to flow and storage constraints (USACE, 1994a). Mathematically, it represents a class of network flow programming, commonly used to solve large logistical, warehousing, and transportation problems in the public and private sectors.

HEC-PRM has been employed to examine reservoir operations in several problem contexts:

- 1) to examine long-term operating strategies for reservoir systems undergoing periodic re-examination (USACE, 1992, 1994b, 1995; Lund and Ferreira, 1996),

- 2) as a screening tool to examine the impact of major changes in system configuration or operating purposes, allowing much more flexible and rapid examination of operating performance than would be required using traditional simulation studies (USACE, 1993a),
- 3) as a tool for illustrating the trade-offs between various operating purposes represented by parties in negotiations over reservoir operation (Kirby, 1994), and
- 4) to suggest promising operating approaches for more real-time and seasonal operations (USACE, 1995, 1996).

To date, HEC-PRM has been applied on the Missouri River main stem system, the Columbia River system, Alamo Reservoir in Arizona, the Highland Lakes in Texas, and the Carson-Truckee system in California and Nevada.

Like any optimization model, HEC-PRM must simplify the reservoir operation problem into a form that can be solved by a computer. For HEC-PRM, some important simplifications include:

- all inflows are known with certainty during the modeled period (This gives the model the opportunity to look ahead to optimize current releases knowing future inflows.);
- evaporation and seepage must be approximated by linear relationships;
- the objective function must be the sum of linear penalty functions;
- hydropower penalties must be approximated by piece-wise linear relationships;
- an individual constraint cannot include flows or storages on more than one arc; and
- all operating objectives must be included quantitatively to be considered.

While few reservoir operations problems will conform exactly to these simplifications, the model still retains an ability to represent many of the most important aspects of most reservoir problems and should be able to provide at least qualitative insight into the problem of optimizing reservoir operations. Indeed, for large reservoir systems, there is no alternative optimization methodology presently available which would not require similar or greater simplifications. The importance of these assumptions can be tested through subsequent simulation studies.

Some caveats for interpreting HEC-PRM results are:

- alternative, equally optimal operating solutions can exist, (although no operating solution will be better),
- model results will be sensitive to the penalty functions used,
- perfect foresight of system inflows is assumed for the analysis period, and
- some of the simplifications mentioned above might be important.

More detailed discussions of HEC-PRM appear in several Hydrologic Engineering Center reports (U.S. Army Corps of Engineers, 1991a, 1994a).

1.7 Implicit Stochastic Optimization

The development of reservoir operation plans by abstracting operating rules from extensive deterministic optimization results is sometimes known as "implicit stochastic optimization" (Whitlatch and Bhaskar, 1978; Klemes, 1979; Karamouz, et al., 1992). The approach relies on using a long record of historical or synthetic hydrologic inflows to represent the uncertainty in inflows. The patterns seen in the deterministic optimization results, which have perfect knowledge of future inflows, should therefore represent optimal rules for operations even under uncertainty.

The major advantages of implicit stochastic optimization over explicit stochastic optimization, such as stochastic dynamic programming and stochastic linear programming, are the much greater computational feasibility of deterministic optimization (Young, 1966) and the relative

ease of establishing input data sets needed for implementing deterministic optimization. Explicitly stochastic optimization methods, for example, typically require an explicit stochastic model of streamflows, which is often elusive. There is even some work to suggest that the rules produced by implicit stochastic optimization are superior to those produced by explicit stochastic optimization under some circumstances, given the different simplifications required for each approach (Karamouz and Houck, 1987).

Ideally, if a deterministic reservoir optimization is performed with a long enough hydrologic record, a contingency table could be developed to establish the mean optimal release from each reservoir given the current month, current storages, and current inflows throughout the system. This was originally done by Young (1966) for a single idealized reservoir using 5,000 periods of synthetic inflows with one season. It is unlikely that this ideal contingency table approach could be developed for most real reservoir systems that have significant monthly variation, multiple reservoirs, and less than a century of hydrologic record.

Nevertheless, implicit stochastic optimization approaches that have lesser requirements and produce more approximate rules have been common in the reservoir optimization literature (Young, 1966; Jettmar and Young, 1975; Whitlatch and Bhaskar, 1978; Bhaskar and Whitlatch, 1980; Trott, 1979; Karamouz and Houck, 1982; Karamouz, et al., 1992). Most applications of implicit stochastic optimization have been to cases with only a short streamflow record, typically less than 40 years. In these cases, use of the historic record would provide only a very limited and perhaps unrepresentative example of the range of streamflow experiences which are possible in the future. In these cases, synthetic streamflow generation has been employed to provide the statistical equivalent of a long streamflow record (Karamouz, et al., 1992). While synthetic streamflow generation may be unavoidable in the absence of a long streamflow record, there are important methodological difficulties with this approach (Klemes, 1974). Still, some have found that the use of even rather short (64-year) historic records can yield operating rules essentially the same as those found using longer synthetic streamflow records (Jettmar and Young, 1975). This approach of using long historical streamflow records as an input to a long-term deterministic optimization model was employed successfully with the Missouri River System (USACE, 1991a, 1994b; Lund and Ferreira, 1996).

1.8 Overview of Report

The next chapter of this report reviews operating rules that can be used for both seasonal and long-term simulation modeling of reservoir systems. Chapter 3 then reviews approaches for developing reservoir operating rules from long-term deterministic optimization results. Chapter 4 reviews the application of deterministic optimization models, such as HEC-PRM, for seasonal reservoir operation studies. Conclusions to this work appear in Chapter 5.

Chapter 2

Seasonal and Long-term Reservoir Operating Rules

2.1 Introduction

All reservoir system simulation models require the specification of reservoir operating rules. These rules determine the release and storage decisions for each reservoir at each time-step during the simulation. This chapter reviews a wide variety of proposed and common operating rules for seasonal and long-term studies. Real-time studies, with an hourly or daily time-step, often have more detailed safety, habitat, and facility limitations not usually important for studies using coarser time-steps or longer operating horizons.

This presentation is divided into four main sections, addressing operation rules for: 1) single reservoirs, 2) reservoirs in series, 3) reservoirs in parallel, and 4) miscellaneous multi-reservoir systems. Each form of reservoir operating rule is presented with a brief description, followed by its applicability and conditions of optimality, examples of its application, and pseudo-code for its implementation in simulation modeling. The applicability of these rules for different operating purposes is summarized in Table 2.1. While this review is not exhaustive, it should be useful for those seeking to develop simulation models or operating rules for reservoir system operations. Other reviews of reservoir operating rules include Sheer's mercifully brief review (1986) and Loucks and Sigvaldason (1982). For those developing operating policies for real reservoir systems, the following rules are a bag of tricks with some theoretical or practical basis to commend them, either individually or mixed together.

Table 2.1: Typically Desirable Single-Purpose and Common Multi-Purpose Reservoir Operating Rules (by section number)

Operating Purpose	Single Reservoir	Reservoirs in Series	Reservoirs in Parallel	General Multi-Reservoir
Downstream Purposes:				
Flood Control	2.2.5	2.3.2	2.4.5	2.5
Navigation	2.2.1-2.2.3	2.3.1	2.4.1- 2.4.4	2.5
Environmental	2.2.1-2.2.3	2.3.1	2.4.1- 2.4.4	2.5
Recreation	2.2.1-2.2.3	2.3.1	2.4.1- 2.4.4	2.5
Water Supply	2.2.2-2.2.4	2.3.1	2.4.1- 2.4.4	2.5
Reservoir Purposes:				
Navigation	2.2.1	-	-	2.5
Environmental	2.2.1	-	-	2.5
Recreation	2.2.1	-	-	2.5
Hydropower Production	2.2.1, 2.2.4,	2.3.3, 2.3.5	-	2.5
Energy Storage	2.2.1	2.3.3, 2.3.4	2.4.1- 2.4.4	2.5
Multiple Purposes:				
	2.2.6-2.2.8	-	-	2.5, 2.6

2.2 Single Reservoir Rules

This collection of operating rules applies to reservoir systems consisting of a single reservoir or reservoirs that are operated independently of one another. (Often in multi-reservoir systems, reservoirs under different ownership are operated with little consideration of other reservoirs.) The eight sections below discuss various single reservoir rules, developed and applied for different operating purposes and conditions.

2.2.1 Fixed Release and Fixed Storage Rules

Often, operating a reservoir to meet a fixed storage or release target best meets the reservoir's operating purposes. These targets can vary seasonally. Fixed storage operations are commonly desirable for reservoir recreation, maximizing head for run-of-river hydropower production, maintaining navigation in the upper reaches of a reservoir, or maintaining fish or wildlife habitat within a reservoir. Fixed release operations are common for providing downstream water supplies or maintenance of downstream navigation, instream environmental flows, or recreation activities.

Applicability and Optimality

Fixed storage or release rules are optimal for several reservoir purposes. For recreation, fixed storage levels keep the water surface high enough for boating or swimming, but not so high as to make docks or parking lots unusable. Fixed releases downstream of reservoirs maintain adequate depths for rafters and boaters. Fixed releases provide minimum depths necessary for pursuing these activities while ensuring that flows are not so swift as to be unsafe or erode channels.

Fixing storage at a reservoir's capacity can maximize hydropower production by maximizing head. This is only feasible if either the reservoir's sole purpose is hydropower production, given sufficient flow, or when other requirements are still satisfied when storage is kept at its maximum level.

Environmental considerations also may warrant fixed storage or releases. Certain habitats or species rely on near-constant water levels within and downstream of reservoirs for growth or survival. While fixed storage rules accommodate the environment within a reservoir, fixed releases support downstream habitats.

Lastly, optimization model results, such as those of HEC-PRM, often imply that relatively constant storage or release levels are desirable for a multi-purpose system. This was found to be the case for several reservoirs on the Missouri and Columbia river systems (USACE, 1992, 1994b, 1995).

Application

The Cumberland River is a tributary of the Ohio River, which runs in a general east-to-west direction across the Kentucky-Tennessee border (USACE, 1985). The primary functions of the Cumberland system are flood control, navigation, and hydropower. For reservoir recreation, efforts are made to maintain the storage reservoirs as high as practicable during the summer months, within constraints of power requirements. At J. Percy Priest and Cordell Hull reservoirs, the authorizing legislation specifies that summer pool elevations be maintained at fixed levels for recreation.

Another project in the Cumberland River system, the Barkley Lock and Dam is operated as a run-of-river project for power production. Because the Barkley Lock and Dam is a small to medium-sized single-purpose power project, it is not suited for providing seasonal storage, so its operation concentrates on hydropower production by maintaining high fixed storage levels.

Modeling and Implementation

Without violating other constraints on a reservoir system, fixed release rules imply merely setting releases to a fixed release target level, D_t , unless inflows exceed available storage (when spill occurs), or if available water is insufficient to meet the release target (shortage cannot be avoided).

$$\text{If } (S_{t-1} + I_t \geq D_t) \text{ then } R_t = \text{Max}(D_t, S_{t-1} + I_t - C) \\ \text{else } R_t = D_t.$$

where

R_t	release decision for period t
D_t	release target for period t
S_{t-1}	beginning of period storage,
I_t	net period inflows (including reservoir evaporation and seepage), and
C	reservoir storage capacity.

Fixed storage is maintained by setting releases for each time period t , R_t , to achieve the storage target, S_{Tt} . This release is given by the equation below:

$$R_t = \text{Max}(S_{t-1} + I_t - S_{Tt}, 0).$$

Comments

Fixed release and storage rules apply to very specific situations. Release and storage decisions become more complex by incorporating other reservoir purposes or when within-reservoir operating purposes conflict with downstream operating purposes, as could be the case for recreational purposes within and downstream of a reservoir.

2.2.2 Standard Operating Policy

For a single reservoir providing water supply downstream, the most straightforward operating policy is to provide as much of the target demand as possible. If insufficient water is available to meet target demands, the reservoir is emptied to make releases as close to the target as possible. Water available in excess of the target release is stored. Water exceeding both the release target and available storage capacity is spilled downstream (Shih and ReVelle, 1994, 1995). The release curve for this standard operating policy (SOP) is shown in Figure 2.1.

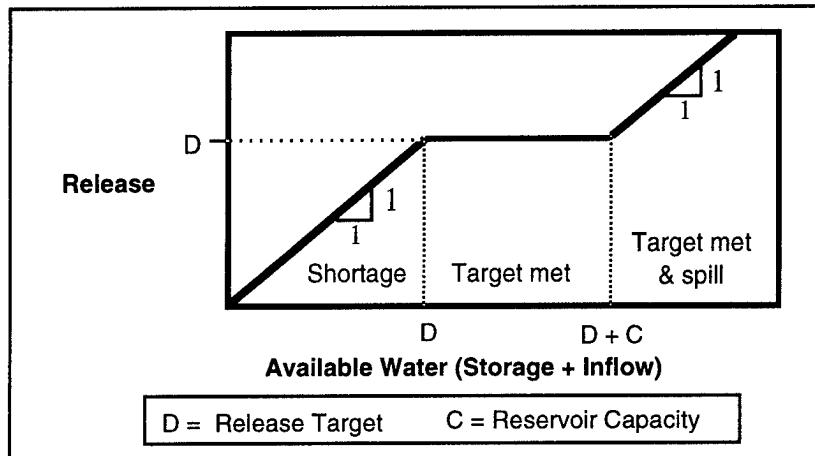


Figure 2.1. Standard Operating Policy

Applicability and Optimality

The optimality of the SOP for water supply purposes depends on both predicted inflows and the economic effect of release shortages (releases below the target). An optimal policy converges to the SOP as either hydrologic or economic uncertainty grows (Klemes, 1977). This rule is also optimal where the relation between economic loss and shortage is both linear and constant between time-steps. This means that the loss due to 5 units of water shortage for one period plus 5 units of water shortage in a second period is economically equivalent to 10 units of water shortage for a single period. When this condition holds, the economic losses from shortages must be linear or, equivalently, the objective is to minimize the average expected value of shortage (Hashimoto, et al., 1982). Where streamflows are highly uncertain (particularly with low or negative temporal autocorrelation), shortages should be postponed for as long as possible in the event that higher than anticipated inflows occur in later time periods (Bower et al., 1966). The SOP maintains this preference of present over future supply when handling shortages.

Application

The so called “standard operating policy” is little used for actual reservoir operations, but is commonly used for operational studies and planning, particularly for firm yield studies.

Modeling and Implementation

The pseudo-code for implementing the SOP can be expressed as:

```

If  $(S_t + I_t) < D_t$  then  $R_t = (S_t + I_t)$ 
else if  $[(S_t + I_t) < (C + D_t)]$  then  $R_t = D_t$ 
else  $R_t = S_t + I_t - C$ 
end if

```

where

S_t storage at beginning of period t ,
 I_t inflow in period t ,
 D_t water demand in period t ,
 R_t release in period t , and
 C reservoir capacity.

Comments

The SOP provides a simple method of making release decisions, although in actual operation, strict implementation of this rule is rare due to a desire to maintain at least some water to avoid extremely severe shortages. When the amount of available water is less than the target demand, rarely would operators completely empty a reservoir. Most reservoir systems employ some sort of water conservation to reduce demands before a reservoir would run dry. The next section discusses "hedging" rules to make release decisions when anticipating that target demands might not be met.

2.2.3 Hedging Rules

Hedging rules allow releases to be curtailed below target levels to avoid potentially more severe shortages later. Hedging can be viewed as a modification of the standard operating policy as depicted in Figure 2.2.

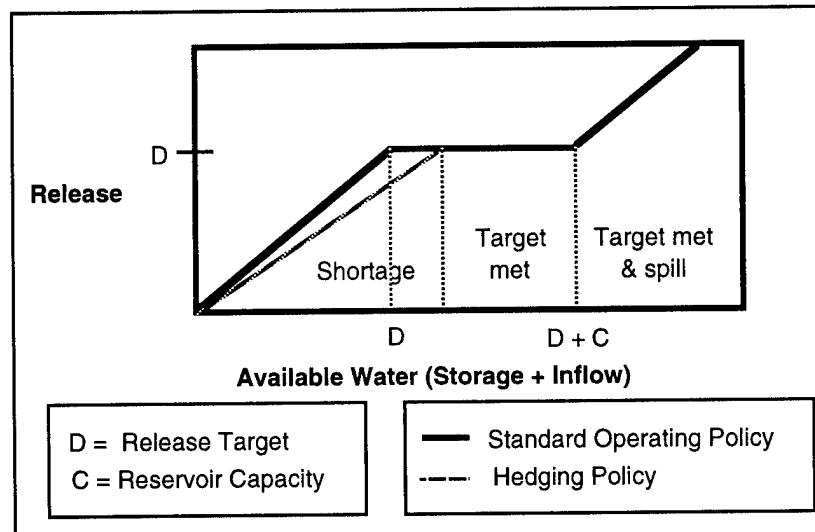


Figure 2.2: Hedging Release Rule

Rather than releasing as close to target releases as possible, as in the standard operating policy, hedging rules accept small deficits in releases when the amount of available water is within a certain range.

Applicability and Optimality

Hedging rules apply to systems that are not always expected to meet target releases and where the severity of shortages is more important than their frequency. Economically, hedging rules should be utilized only when losses are not linearly related to shortages (Bower et al., 1966). As discussed previously, the standard operating policy yields optimal results when losses are linearly related to shortages. However, where increased shortages incur disproportionately greater costs, it is often wise to reduce the frequency of experiencing large shortages by increasing the

frequency of small shortages. If doubling the shortage results in *more* than twice the original loss, then dividing the total expected shortfall into more frequent but less severe shortages can minimize total losses.

The applicability of hedging rules also depends on the nature of a watershed's hydrology. If droughts are isolated in time and of short duration, then creating a small shortage to maintain some reservoir storage might not provide additional useful water. Here, hedging releases, and creating small shortages, when later reservoir inflows prove to be more than sufficient is futile. Such futile hedging incurs the costs of small shortages without a compensating reduction of losses from large shortages. However, where there is substantial persistence in low-inflows, the total shortage often can be more evenly distributed using a hedging rule.

Application of the hedging rule depends on both predicted inflows and the economic effect of release shortages. If enough hydrologic or economic uncertainty exists, hedging can be no longer optimal, and the standard operating policy becomes a sufficient operating scheme (Klemes, 1977).

Incorporation of Performance Measures

Determination of the "optimal" hedging rule requires some measurement of system performance. Often, this measurement of system performance is economic, with an objective of minimizing economic losses. However, other less economically based, performance measures are commonly suggested such as reliability, resiliency, and vulnerability. Reliability refers to the probability that the reservoir can supply a target release level. Resiliency describes how quickly a system returns to a satisfactory state once a failure to provide target releases has occurred. And vulnerability is the likely magnitude of a failure, if it occurs (Hashimoto, et al., 1982).

Some relationships exist between these performance measures that may be useful in determining the best hedging policy. Operations of a reservoir that are both reliable and resilient reservoirs can be vulnerable to failures of large magnitude (Moy, et al., 1986). Reductions in the rate of failure tend to increase the likelihood of recovery (Hashimoto, et al., 1982).

In another study, several relationships were observed between the performance measures and mean deficits by varying parameters P1 and P2 as shown in Figure 2.3 (Bayazit-Unal, 1990). Here, P1 is the minimum amount of available water at which hedging is applied (the left most point of the dashed line), while P2 is the amount of available water above which hedging is no longer applied (the right most point of the dashed line). For the hydrologies examined:

1. Hedging reduces the mean deficit and vulnerability when P1 is close to the target volume and P2 is large (near the target volume plus reservoir capacity). This means that hedging is continued long after the full demand can be met.
2. Reliability and resiliency are not affected by the value of P1 but are reduced significantly with increasing values of P2.
3. The standard deviation of all performance criteria including mean deficit are reduced when P2 is increased.
4. If hedging is applied when there is little water in storage, the mean deficit and vulnerability will be increased, while the standard deviation of vulnerability will be significantly reduced. This may be desirable in preventing catastrophic future deficits.
5. The standard operating policy provides the best levels of reliability and resiliency with satisfactory deficits.

The tradeoffs that exist among the performance measures may be important in establishing the optimal hedging policy (Bayazit-Unal, 1990).

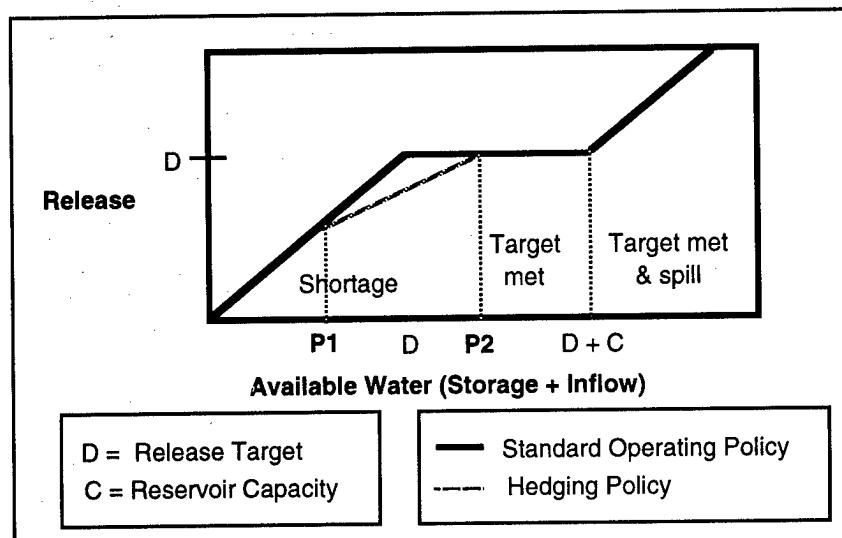


Figure 2.3: P1-P2 Hedging Rules

Applications

Most reservoir systems employ some sort of water conservation when storage levels become low before the reservoir runs dry. In the Occoquan Reservoir in Virginia, hedging is implemented when the amount of available water reaches a critical level. When water levels drop below this level, specific emergency measures are invoked. These measures involve the purchase of water from other sources and various degrees of hedging. A conservation factor, α , determines the level at which customers are restricted in their water usage. This factor ranges from 1.0 to 0.8, where a value of 1.0 indicates that customers are unrestricted. Table 2.2 shows the various levels of hedging corresponding to the amount of available water in storage.

For storage volumes above 4,300 Mgall, no hedging is required. Rather than having levels of conservation or hedging for every possible level of storage, the operation has been simplified by the discretization of storage levels (Hirsch, 1978).

Table 2.2. Hedging in Occoquan Reservoir Operations

Volume in storage at beginning of the month (Mgall)	Emergency stage invoked for the month	Anticipated Conservation, α
9,800 - 4,300	0	1.00
4,300 - 3,400	I	0.95
3,400 - 1,900	II-A	0.90
1,900 - 1,100	II-B	0.80
1,100 - 0	III	0.80*

*No estimate is made of the amount of conservation in this stage. The economic consequences of this stage are severe. When it occurs the local government will require the curtailment of all water uses not essential to health and safety.

Modeling and Implementation

Several forms of hedging are available for modeling and implementation. Zone-based hedging is the most commonly used form of hedging (See Zone-based operations). The implementation of hedging for the Occoquan Reservoir is an example of zone-based hedging,

where the releases are specified as fractions of normal demand rather than specific release values. Five forms of hedging are described below.

Trigger Hedging Value (One Point Hedging)

Hedging can be implemented by specifying a trigger value, the amount of available water at which hedging begins, and the degree of hedging below this volume. This is accomplished by choosing a single constant K_t which may be unique for each period t . The trigger volume is set to be K_t times the period's demand or release target, D_t . When the amount of available water falls below this volume, the release, R_t , is:

$$R_t = \frac{(S_t + I_t)}{K_t}$$

The storage at the end of the previous period plus the prediction of current inflow makes up the amount of water available. When this amount exceeds the trigger volume (but is less than the demand plus capacity), the full demand can be released for the current period. Figure 2.4 gives a graphical representation of this type of hedging:

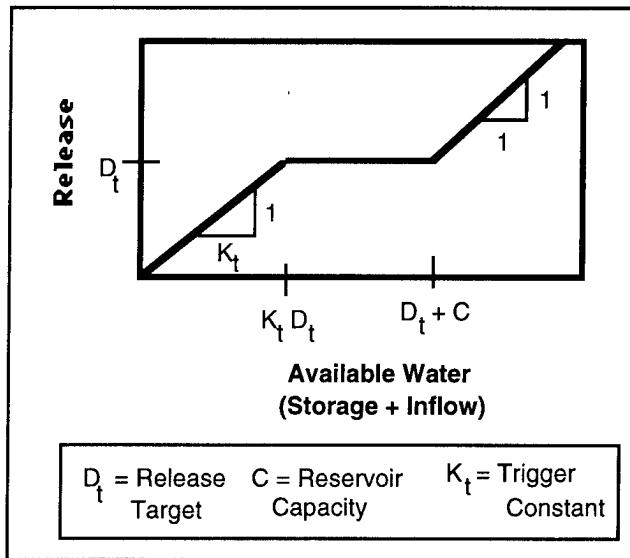


Figure 2.4: Trigger Value (One Point) Hedging

Trigger value or one-point hedging can be implemented using the following pseudo-code:

```
If  $(S_t + I_t) < (K_t D_t)$  then  $R_t = (S_t + I_t)/K_t$ 
else if  $(S_t + I_t) < (C + D_t)$  then  $R_t = D_t$ 
else  $R_t = S_t + I_t - C$ 
```

where

S_t storage at beginning of period t ,
 I_t expected inflow for period t ,
 K_t trigger volume constant for period t ,
 C reservoir capacity,
 D_t demand for period t ,
 R_t release in period t .

For all cases, $K_t \geq 1$. The determination of the K_t value that will minimize losses can be modeled and solved as a mixed-integer mathematical programming problem (Shih and ReVelle, 1994).

Two-Point Hedging (P1 & P2)

Another form of hedging makes use of two parameters, P1 and P2. In Figure 2.3, P1 is the amount of available water at which hedging begins (the left most point of the dashed line) while P2 represents the amount of available water above which hedging is no longer applied (the right most point of the dashed line). Outside this range, the standard operating policy is applied. Values of P1 and P2 can be chosen that will optimize a desired objective, such as minimizing losses (Bayazit and Unal, 1990). The following pseudo-code implements the P1 and P2 hedging rule:

```

If  $(S_t + I_t) < P1$  then  $R_t = S_t + I_t$ 
else if  $(S_t + I_t) < P2$  then
     $R_t = \left[ \left( \frac{D - P1}{P2 - P1} \right) (S_{t-1} + I_t) \right] + P1 - P1 \left( \frac{D - P1}{P2 - P1} \right)$ 
else if  $(S_{t-1} + I_t) < (C+D)$  then  $R_t = D_t$ 
else  $R_t = S_{t-1} + I_t - C$ 

```

where

S_t storage at beginning of period t ,
 I_t expected inflow for period t ,
 C reservoir capacity,
 D_t demand for period t ,
 R_t release in period t .

Three-Point Hedging (P1, P2, and P3)

If an intermediate point, P3, is used in conjunction with the P1 and P2 parameters used in the previous form of the hedging rule, the dashed hedging line in Figure 2.3 would no longer consist of one straight line but two contiguous straight lines. This adds some flexibility to the hedging policy, but also complexity to the operating rule. In the same manner, any number of intermediate variables could be included.

Zone-based Hedging

In zone-based hedging, release is determined by the amount of storage at the beginning of the period. Different levels of hedging are implemented between storage levels A_S , B_S , C_S , by using the following pseudo-code:

```

If  $S_t < A_S$  then  $R_t = A_R$ 
else if  $S_t < B_S$  then  $R_t = B_R$ 
else if  $S_t < C_S$  then  $R_t = C_R$ 
else  $R_t = D_t$ ,

```

where

S_t amount of storage at the beginning of the period t ,
 R_t amount of release for period t ,
 A_S, B_S, C_S specific values of storage, $A_S < B_S < C_S$
 A_R, B_R, C_R specific release values, $A_R < B_R < C_R < D_t$, and
 D_t demand in period t .

Zone based hedging is graphically shown in Figure 2.5.

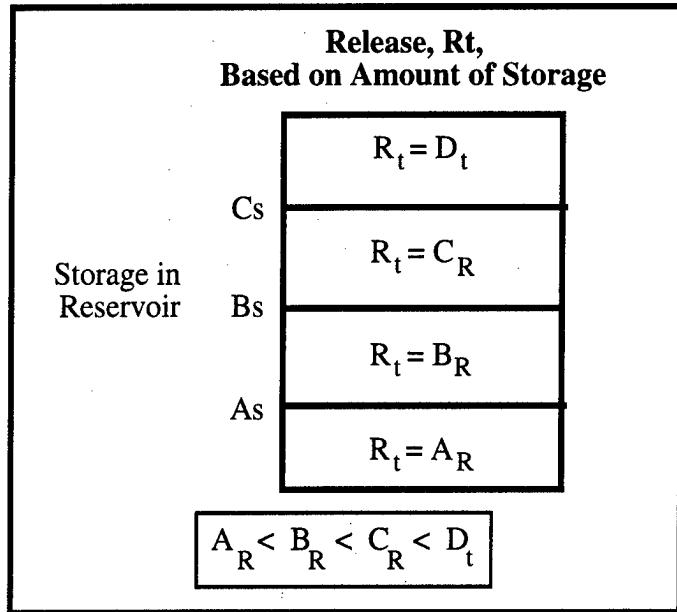


Figure 2.5: Zone Based Hedging

Stochastic Dynamic Programming Hedging

For a given loss function, $L(R)$, and certain explicit probabilistic descriptions of inflow hydrology, stochastic dynamic programming can be used to determine operating policies that will minimize losses due to shortages. If the loss function is convex (exponent, $\beta > 1$), expressed in terms of release, R , and target, D ,

$$L_\beta(R) = [(D - R) / D]^\beta,$$

optimal operating policies exhibit hedging. The parameter β reflects the non-linearity of the loss function. The initiation and extent of hedging applied depends on the value of β and the initial amounts of water available as shown in Figure 2.6 for $\beta = 3$ (Hashimoto et al, 1982). Here, the diagram is modified to include both initial storage S and storage plus inflow ($S+I$) as independent variables, since temporal autocorrelation is included in the stochastic dynamic programming formulation. The lines show the best value of release R as a function of initial storage S plus inflow for specified value of S and release target D ($= 4.5 \times 10^7 \text{ m}^3$).

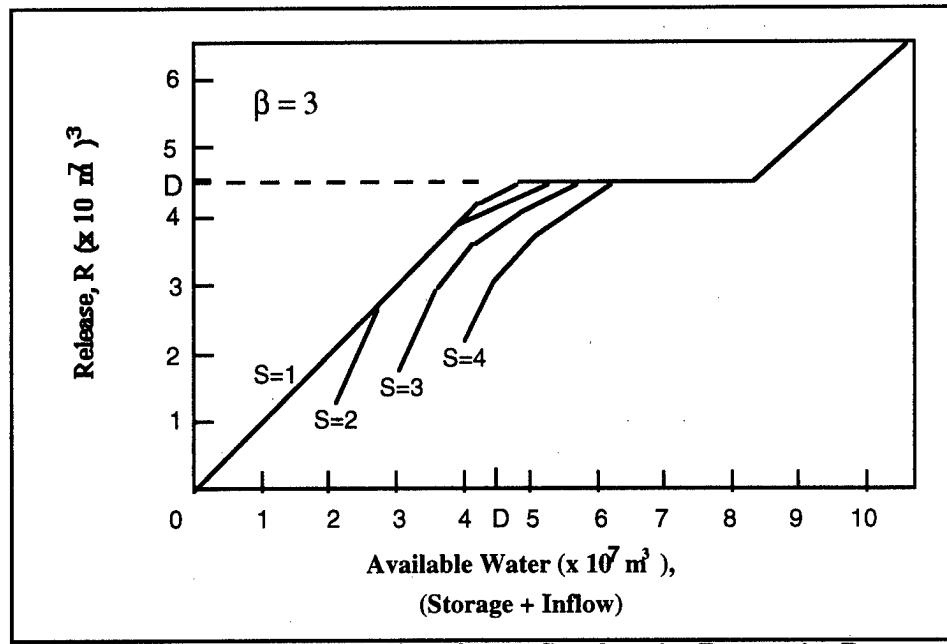


Figure 2.6. Optimal Hedging Policy from Stochastic Dynamic Program

Comments

Hedging is common in the operation of reservoirs for water supply purposes. However, there is often little explicit optimization of hedging operations to maximize system performance over a range of conditions. Thus, for all hedging rules examined, significant efforts must be made to select values of hedging rule parameters (P_1, P_3, P_5 , etc.) which result in desirable operations. This selection of parameters can be accomplished either by simulation or optimization studies (Hashimoto, et al., 1982).

2.2.4 Pack Rules

In contrast to hedging rules which attempt to prepare for expected shortages, pack rules allow beneficial use of water available in excess of targets.

Applicability and Optimality

In single purpose systems pack rules are useful whenever releases above specified targets are of value. Examples include systems with firm and secondary hydropower, water supply, water-pollution control, and ground water recharge.

Bower et al., (1966) discuss application of a pack rule for hydropower production in excess of firm power production targets, where additional power production has some economic value. With the application of the pack rule, end of season spills are avoided by releasing above target levels (but within the hydraulic capacity of the turbines) toward the end of the refill season. In Bower, et al.'s example, hydropower releases are made beyond target levels in anticipation of flows above target release levels in the future. This allows the use of these excess flows to generate secondary hydropower, which has some economic value. Given perfect forecasts, the use of a pack rule should allow the avoidance of spills, unless average flows exceed the turbine release capacity for the period of study or storage capacity is limiting.

A pack rule for water supply purposes might be the release of additional water early in a high-runoff season for pre-irrigation or other lower valued uses of water. For systems with multiple water sources, release of additional water from high-quality or inexpensive sources early in the season can make storage available for such higher-valued waters later in the year, decreasing overall cost and improving the quality of water supplied.

To be applicable, flows in excess of target releases must have some value, as shown in Figure 2.7. If releases above target levels were not useful, then pack rules can not improve system performance. Releasing above target levels would be equivalent to spilling. In cases where there is hydrologic uncertainty, the potential value and likelihood of additional performance resulting from use of a pack rule must be weighted against the cost and likelihood of shortages resulting if inflows fall below predicted levels.

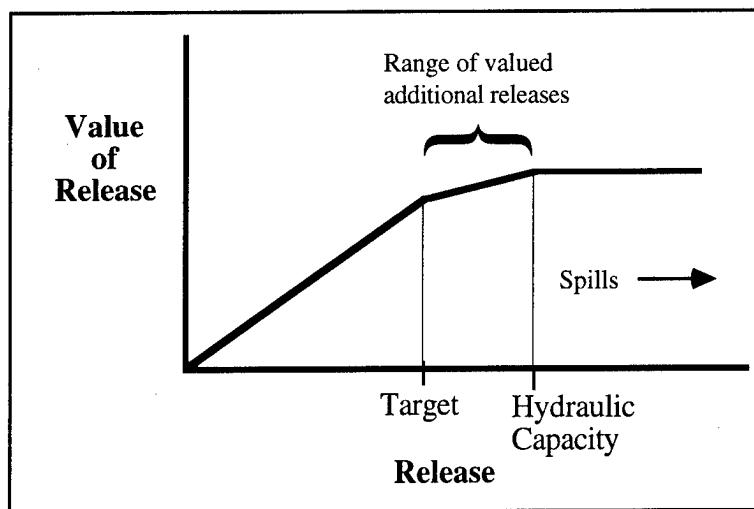


Figure 2.7: Value of Release

Figure 2.7 shows that the increase in value of additional releases below the target is greater than additional releases above the target. If this were not the case then optimal reservoir operation

would consist of large periodic releases. Thus, the value function for releases must be concave for pack rules to be useful. Additional conditions necessary for pack rule implementation are that it applies to an individual single purpose reservoir and is used in the refill season.

Application

Pack rule operations are common in hydropower systems where flows are anticipated in excess of those needed to meet firm contracted power demands.

Modeling and Implementation

The release during any period t is composed of the target release, R_{Tt} , plus the additional release above the target, R_{dt} :

$$R_t = \text{Max}(R_{dt} + R_{Tt}, R_{Tt})$$

$$R_{dt} = Q_{n-t} - (C - S_{Tt}) - P_{n-t}$$

where

R_t	release in period t ,
n	number of periods in drawdown-refill cycle,
t	current period,
Q_{n-t}	predicted flow into reservoir for rest of cycle,
C	reservoir capacity,
S_{Tt}	reservoir storage in current period after current inflows and target release,
and	
P_{n-t}	useful release before spilling for rest of cycle.

If a negative value for R_{dt} is obtained, the target release should just be met. The amount of additional release, R_{dt} , must not be greater than reservoir storage nor exceed the maximum useful release above the target. In energy production, for example, the capacity of the turbines would constitute the amount of useful release.

Comments

Just as hedging rules are used to select operations under real and potential shortage conditions, pack rules are used to select operations where water is available in excess of water release targets. While pack rules operate under less dire circumstances, cases of water availability above target levels are usually more common than drought operations.

2.2.5 Flood Control Rules

This section examines long-term and seasonal flood control rules used for single or independently operated reservoirs. Flood control rules are designed to reduce damages from excessive downstream flows. Typically, flood control rules attempt to minimize peak streamflows at specified downstream locations. However, in some cases, the duration of flooding also is important.

Much of reservoir operation during flood events occurs at an hourly or daily time-step, even though reservoir operations anticipating more flood-prone seasons may extend over several months. These short-term or real-time flood control operations also often must consider issues of keeping rates of change in discharge and reservoir elevation below levels that might endanger downstream users (e.g., recreation) or the geotechnical or structural stability of the reservoir rim, dam, or outlet works. Long-term and seasonal rules representing flood control operations (at a weekly or monthly time-step) are inevitably approximations of actual operations and involve often significant averaging of operations and flows. While this averaging is sometimes acceptable, it is

sometimes valuable to model flood control operations at a smaller time-step (during flood events), when a variable time-step length can be employed in a simulation model. However, modeling flood events at a smaller time-step can introduce additional complications and uncertainties regarding flood-wave routing downstream.

Applicability and Optimality

Several types of flood control rules are available for single reservoirs. Each rule type attempts to maximize the use of storage to reduce flood peaks downstream. These rules tend to apply to either the filling or evacuation periods of flood control operations. Figure 2.8 illustrates several of the rules discussed.

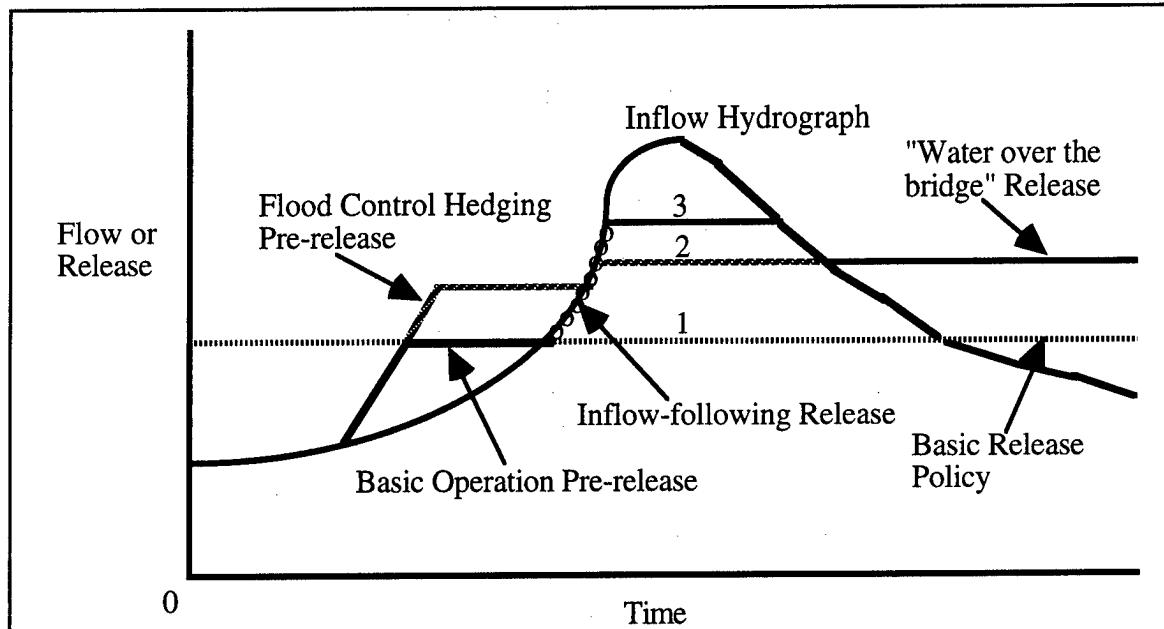


Figure 2.8: Flood Control Storage Fill and Evacuation Rules

Basic Single Purpose Flood Control Rule

For a single reservoir whose sole purpose is to provide flood protection, the downstream flood peak is minimized by keeping a maximum amount of flood storage available, except for the capture and slow release of flood peaks (USACE, 1985). Typically, storage is released from a flood control reservoir at the rate of the downstream channel's maximum flow which incurs little or no flood damage (USACE, 1976). By releasing at or near this critical discharge level (including compensation for downstream inflows), a reservoir has the greatest ability to prevent downstream flooding by freeing storage for the control of later hydrograph peaks. Releasing at this maximum channel capacity also has benefits of maintaining existing channel capacity by clearing the channel and discouraging development in the channel (USACE, 1976). As illustrated in Figure 2.8, pre-releases above current inflows, but below the flood discharge capacity downstream, are sometimes made in anticipation of higher inflows to create additional flood control storage to capturing more the anticipated hydrograph peak.

Whether or not flooding occurs depends on the capacity of the reservoir, the intensity and duration of the inflows, and local inflows below the reservoir. This basic set of rules, while simple, is not necessarily optimal where the future flood hydrograph might exceed the reservoir's capacity and lead to damaging flows downstream.

Flood Control "Hedging" Pre-Release Rules

When it is expected that flooding can not be avoided completely, it may be wise to employ a type of flood control "hedging". In this operation, small amounts of flooding are allowed to occur downstream to avoid potentially more devastating flooding later. This occurs by making pre-releases exceeding the normal downstream channel capacity, as illustrated in Figure 2.8. As with water supply hedging, this type of flood control should only be used when the occurrence of flooding is expected with a high degree of certainty and when a single moderate flood is more destructive or costly than the total of a series of smaller flooding events. This approach to flood control operations increases the ability to respond to moderate floods by maintaining some control over reservoir releases and spill.

Inflow-Following Filling Rules

A common flood control filling rule is to not release more than the inflow, if inflows exceed the flood discharge capacity downstream, as shown in Figure 2.8. In the event that anticipated inflows fall below expectations, this rule minimizes the likelihood that operating authorities will increase flood damages by their operation of the reservoir. This rule may often be necessary to minimize operating liability for downstream flood damage. However, for moderate to large flood events (relative to the reservoir's storage capacity and downstream channel capacity), this policy may preclude pre-releases needed to minimize downstream flood damage.

The resulting peak flows from the various flood storage filling rules also are illustrated in Figure 2.8. For small floods (for given storage capacity and inflows) the basic flood control operating rule is sufficient, with peak flow (1) falling at or below downstream channel capacity. The basic flood control operating rule can be extended to being sufficient for some moderate floods if pre-releases (less than downstream channel capacity) can be made to expand the flood storage capacity available (1, with basic operation pre-release). Setting releases to follow inflows exceeding the downstream channel capacity for a time before storage is filled allows for best capture of the flood peak for moderate and large floods (3). However, if pre-releases can be raised above the downstream channel capacity to further increase available flood control storage, more of the peak can be captured and peak flows might be further reduced (2).

"Water over the Bridge" Storage Evacuation Rules

If circumstances make it impossible to prevent flooding and flood damage has already occurred downstream, it may be beneficial to make releases that result in downstream flows greater than the channel capacity, up to flows equaling the greatest level of recent inundation. The inundated areas may be able to handle the additional releases without substantial additional damage, allowing the creation of additional flood control storage in the event that inflows become more severe in the future. This type of operation might not be desirable if prolonged flooding duration greatly increased damage in already flooded areas, such with the destruction of some crops, orchards, or habitats which might survive brief flood peaks, but suffer from prolonged flooding (USACE, 1979b). In real time operations, downstream inflows can make this release strategy difficult without intensive monitoring and forecasting of downstream inflows. However, this approach can also accelerate the rate at which reservoir storage becomes available for subsequent floods.

Zone-based Flood Control Rules

Flood control is more commonly part of a multi-purpose reservoir system where non-flood control benefits are achieved when storage levels are kept near maximum rather than minimum levels. Hydropower production, water supply, and recreation typically all benefit from keeping storage levels high throughout most of the year.

Zone-based flood control rules reflect a compromise between flood control and other operating purposes by dividing the reservoir into several zones or pools, each with a different storage capacity and operating rule. The size and operation of each zone can also vary seasonally. Flood control zones typically occupy the top zones, as shown in Figure 2.9.

In reservoirs providing flood protection, the top pool is often reserved for exclusive flood control. Water encroaching into this zone is evacuated as quickly as possible without exceeding downstream channel capacities or other maximum release restrictions. Flood control operations also are reflected in the operating rules for the second highest pool, the flood control and multiple use zone in Figure 2.9. The amount of water in this zone is reduced to a prescribed level by a certain time to reserve space for the capture of high inflows. The size and timing of the seasonal flood control zone can be made to vary with watershed conditions. During dry years, when there is little moisture in the watershed, the flood control operation of a flood control and multiple use zone may begin later and be reduced in size, while in wet years the flood control operation of this zone may begin earlier and make up a greater portion of overall capacity. Such variations in the operation of non-exclusive flood control zones are typically explored using long-term studies (usually to set rules or zones), with decisions made annually through annual, seasonal, or real time operation studies.

Storage can be kept at a range of levels as long as the rules defined for each zone are observed. The next section in this chapter (Section 2.2.6) discusses the functions and operations of the lower two zones. In essence, zone-based reservoir operations allow the creation of vertically separated reservoirs, each with different release rules, typically for different purposes.

Mariën, et al. (1994) present both theory and methods for establishing flood control storage targets for single reservoirs and individual reservoirs within a system. These rules are built around the mathematical "controllability conditions" necessary to avoid exceeding flood flows at particular points within the system.

Application

Flood control plays an integral role in the operation of the Missouri River system. The main stem of this system spans seven mid-western states and is composed of six reservoirs followed by a long stretch of river between Sioux City, Iowa, and the confluence with the Mississippi River at St. Louis, Missouri. Each main stem reservoir is partitioned into four operational zones, two of which involve flood control operation, as in Figure 2.9 (USACE, 1979a).

The top exclusive flood control pool is evacuated as soon as possible without exceeding the safe bankfull capacity of downstream channels. The exclusive flood control zone and the second highest pool, the flood control and multi-purpose zone, must be completely evacuated by March 1 of each year to make available storage for high spring and summer inflows. The lower two zones are not directly involved in flood control.

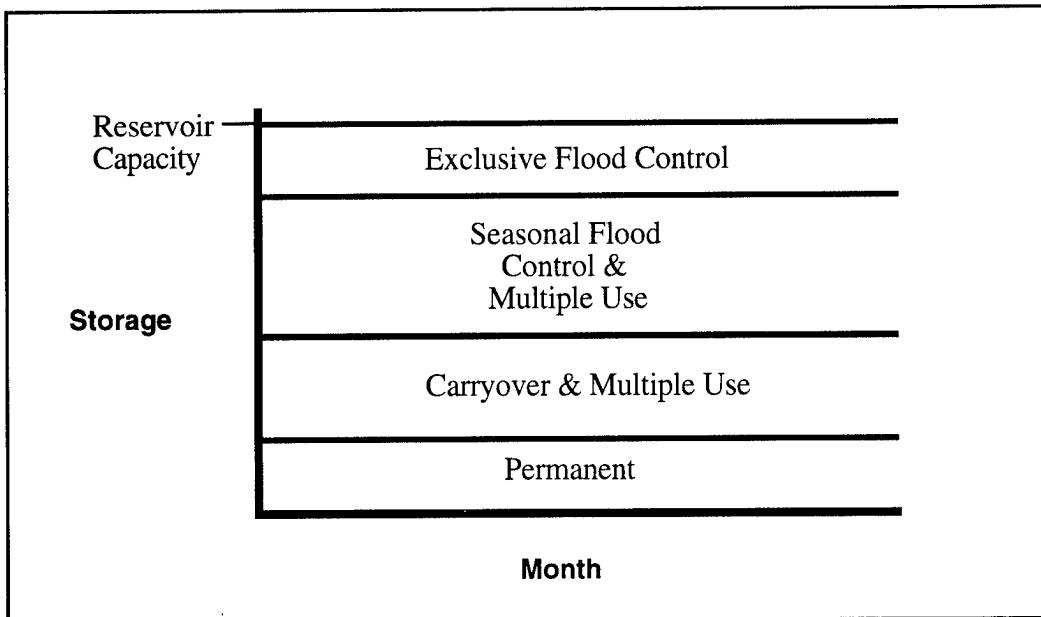


Figure 2.9: Missouri River System Zone Based Operation

Modeling and Implementation

Each of the foregoing flood control operating rules imply different algorithms for simulation modeling.

Basic Single Purpose Flood Control Reservoir Rule

For a single reservoir used exclusively for flood control, the following pseudo-code can be applied:

```

If  $[(S_t + I_t) \leq Q_{c,t} - f_t]$  then  $R_t = S_t + I_t$ 
elseif  $[(S_t + I_t) \leq (Q_{c,t} - f_t + K)]$  then  $R_t = Q_{c,t} - f_t$ 
else  $R_t = (S_t + I_t) - K$ 

```

where

S_t storage at beginning of period t ,
 I_t expected inflow for period t ,
 $Q_{cr,t}$ downstream channel capacity for period t
 f_t downstream (local) inflows above the flood control location,
 R_t release in period t , and
 K reservoir capacity.

This rule has a form similar to the Standard Operating Policy for water supply, replacing the target release with the critical flood discharge downstream (corrected for downstream inflows). Total water availability (inflow plus storage) exceeding reservoir storage capacity plus the critical flood discharge results in flooding.

Flood Control "Hedging" Pre-Release Rules

Flood control hedging rules would essentially increase $Q_{c,t}$ as the reservoir filled and the prospects for high future inflows increased. This increase in $Q_{c,t}$ would result in some flooding downstream, but preserve some flood storage to allow some control over potentially worsening conditions. Small floods would become more frequent, while the frequency of moderate floods would be reduced.

Inflow-Following Filling Rules

Inflow-following flood storage filling sets releases equal to inflows, retaining flood control storage. When a desired target downstream release (Q_T) is exceeded, available storage capacity is filled to capture the hydrograph peak. This operation is given by the following pseudo-code:

```
If ( $I_t \geq Q_T$ ) then  $R_t = Q_T$ 
elseif ( $I_t \geq Q_{c,t} - f_t$ ) then  $R_t = I_t$ 
else other rule for emptying flood control storage.
```

"Water over the Bridge" Storage Evacuation Rules

The "water over the bridge" rule modifies the simple flood control rule merely by increasing the allowable downstream flow to reflect damage already incurred from earlier flooding. For a single reservoir used exclusively for flood control, the following pseudo-code can be applied:

```
 $Q_{pt} = \text{Max}(Q_{\tau=0}, \dots, Q_{\tau=t-1}, Q_{c,t})$ 
If [ $(S_t + I_t) \leq Q_{pt} - f_t$ ] then  $R_t = S_t + I_t$ 
elseif [ $(S_t + I_t) \leq Q_{pt} - f_t + K$ ] then  $R_t = Q_{pt} - f_t$ 
else  $R_t = (S_t + I_t) - K$ 
```

where Q_{pt} is the peak flow downstream to date during flood control season. The first line established Q_{pt} as the larger of the greatest previous flow of the season or the flood channel capacity downstream.

Zone-based Flood Control

In zone-based operation, the exclusive flood control zone is kept empty by implementing the following pseudo-code:

```
 $F_t = K - V_{FCt}$ 
If ( $S_t + I_t > F_t$ ) then
  If ( $S_t + I_t - F_t \geq Q_t$ ) then  $R_t = Q_t$ 
  else  $R_t = (S_t + I_t - F_t)$ 
else {release depends on other priorities during current period}
```

where

F_t	reservoir volume not devoted to exclusive flood control in period t (including storage allocated to all other zones),
K	reservoir capacity,
V_{FCt}	storage capacity reserved for flood control in period t ,
S_t	storage at beginning of period t ,
I_t	expected inflow for period t ,
Q_t	critical flood discharge for period t , and
R_t	release in period t .

Comments

The periods used in flood control release decisions are commonly reduced from the usual monthly time steps to days and even hours due to storms and flash flooding. Flood control operations are thus among the most delicate and difficult operating decisions, they must be made quickly, with significant and immediate consequences, and typically must be made in situations with considerable future uncertainty, and often considerable uncertainty regarding the present condition of the system.

2.2.6 Zone-Based Operations

In zone-based operations, individual reservoirs are divided into virtual storage zones or pools (USACE, 1977; Loucks, et al., 1995). Operating decisions are then based on the amount of water contained within a reservoir in addition to other decision parameters. This method of reservoir operation was first discussed in terms of Flood Control Rules and is graphically shown again in Figure 2.10. Operating decisions depend on the zone corresponding to the current reservoir storage level. The size of each zone may vary from month to month.

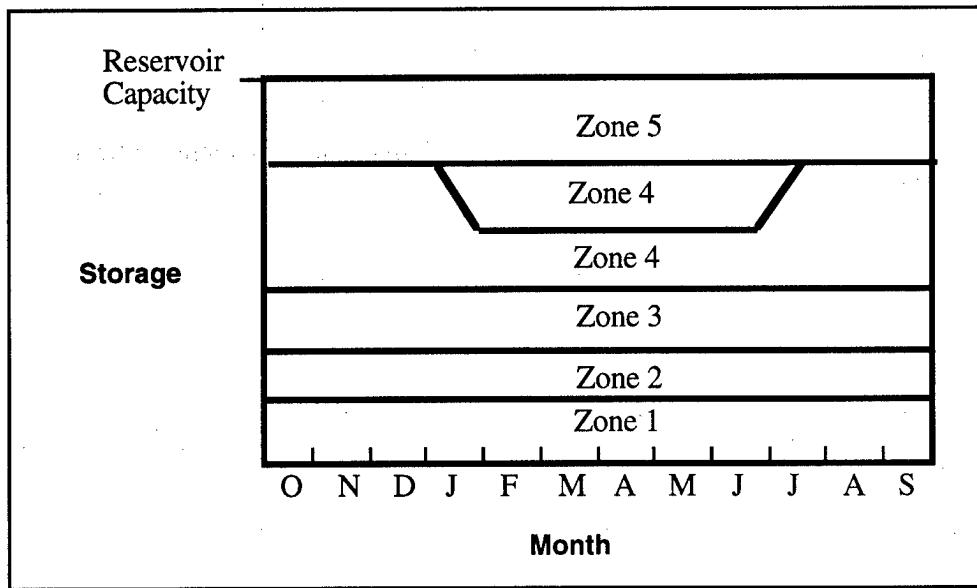


Figure 2.10: Zone-Based Operation

Applicability and Optimality

Single-Purpose Zone-based Operation

For single purpose reservoirs, each zone typically has a target release. The use of zones is well suited to single purpose reservoirs providing water supply, for example. When storage levels become low and water conservation becomes necessary, the reduction of releases is determined by which zone current storage resides. This type of operation was presented earlier as zone-based hedging.

Multi-Purpose Zone-based Operation

When reservoirs serve multiple purposes it becomes difficult to develop simple rules to define the operating policy. Dividing the storage capacity into zones provides a mechanism to deal with the various purposes of the reservoir. Typically, each zone has a specific function and rules associated with it. Each zone may be assigned a particular function, such as flood control, hydropower, navigation, water supply, recreation, or low-flow augmentation. One zone may also define the range of storage utilized in a typical annual drawdown-refill cycle. The number and function of the zones varies with the functions of the reservoir system as well as with time but can consist of a flood control zone, conservation or power storage zone, buffer or annual carryover zone, and a permanent zone. Some commonly used zones are discussed below, as they appear in Figure 2.10.

Exclusive Flood Control (Zone 5). The top zone is reserved for exclusive flood control. It is operated to capture flood waters, but is to be evacuated as soon as possible without exceeding maximum release levels. The boundaries of this zone are the capacity of the reservoir and the highest storage target attained during a typical annual drawdown-refill cycle.

Flood Control and Multiple Use (Zone 4). This zone provides initial flood storage before use of the exclusive flood zone. The “normal” flood flows that fill this zone are used later for water supply, navigation, hydropower, and other downstream uses. It is within this zone that the annual drawdown-refill cycle is to occur. Flood control operation of at least part of this pool is common seasonally. The flood control and multiple use zones also are termed conservation, water supply, or power storage zones.

Carryover and Multiple Use (Zone 3). Sometimes also referred to as the buffer zone, the carryover and multiple use zone is essentially a drought reserve. It is utilized when the flood control and multiple use zone contains insufficient stored water to meet downstream demands. Typically, this zone is defined by the lowest storage level attained during a typical annual drawdown-refill cycle and the lowest storage attained in a repeat of the drought of record.

Permanent Storage (Zone 2). This storage is intended to remain inactive. It provides a minimum storage level for hydropower, recreation, sediment storage, fish and wildlife, and assures a minimum level for pump diversions.

Dead Storage (Zone 1). Dead storage is the lowest layer of reservoir storage and is part of the permanent or inactive storage zone. At this level the water is inaccessible for withdrawal. The dead storage zone often is used for hydropower or diversion head, to trap sediment, or to reduce the turbidity or temperature of releases or withdrawals.

Application

The Arkansas River drains parts of seven southwestern states and empties into the Mississippi River (USACE, 1985). The operations in this reservoir system are divided into three zones as shown in Figure 2.11. The seasons in which floods and droughts occur could overlap, so it is impractical to provide a joint-use zone to serve for both flood control and conservation storage.

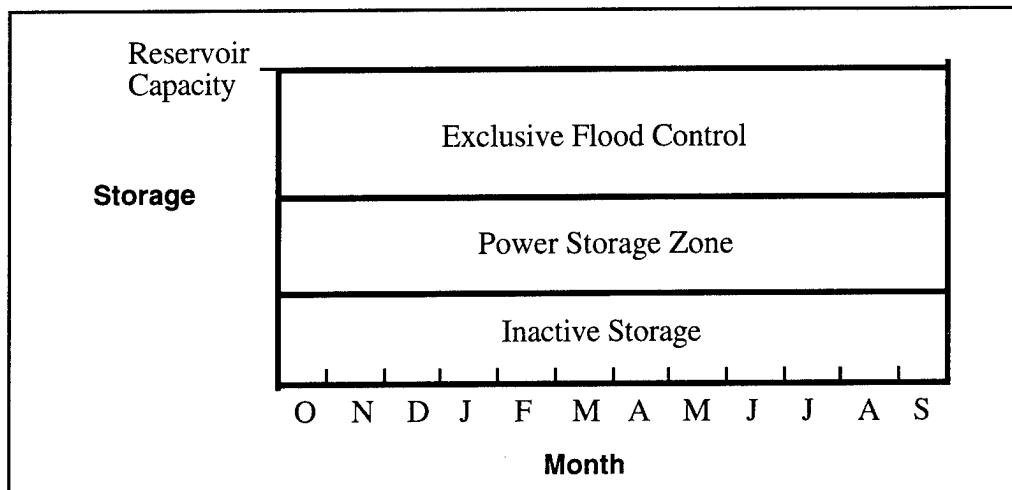


Figure 2.11. Arkansas Zone-based Operation

Normally water entering the exclusive flood control zone is evacuated as soon as possible but this results in spilled energy if discharge exceeds the hydraulic capacities of hydropower plants. Rapid evacuation also causes agricultural damage and navigational problems. Therefore, discharge levels are gradually reduced as flood control space is evacuated. When 40 percent or more of the basin flood control storage is filled, releases are scheduled so as not to exceed 150,000 cfs. When flood control storage is evacuated to 40 percent storage, releases are gradually reduced, so that by the time storage levels are in the 10-16 percent range, releases are at 105,000 cfs, the limit of agricultural damage. Releases are further reduced to 20,000 to 40,000 cfs as flood storage continues to be evacuated.

Within the power storage zone, storage is regulated to meet firm energy and water supply requirements. Firm energy represents the maximum energy a system can provide if the worst historical drought were to reoccur. To completely meet energy demand, frequent drafting would be required which would lower reservoir levels and hydropower generating capacity. The alternative is to keep reservoir levels high enough to protect their dependable capacity and then to make additional thermal energy purchases to meet overall energy demand. The firm energy requirement makes this alternative the practiced operation within the power storage zone (USACE, 1985).

Modeling and Implementation

Single-Purpose Zone-based Operation

The pseudo-codes provided for zone-based hedging and zone-based flood control operations are examples of how these rules can be implemented for single purposes.

Multi-Purpose Zone-based Operation

Multi-purpose zone-based operations implement a specific set of rules varying with the level of current reservoir storage, S_t . Loucks et al. (1995) describe the implementation of zone-based rules for single-reservoir simulation. This can be implemented generically:

```
If  $S_t$  is within Zone 4 then implement Zone 4 rules
  else if  $S_t$  within Zone 3 then implement Zone 3 rules
  else if  $S_t$  within Zone 2 then implement Zone 2 rules
  else implement Zone 1 rules.
```

Comments

Zone-based operations are also used in multi-reservoir systems, but the release decisions remain based on individual reservoir zone rules. These operations provide just one mechanism for handling a complex variety of demands imposed on a reservoir system. Other operating rules that have the capability to handle multi-purpose systems are the complex operating rules discussed in the following sections.

2.2.7 Storage Target Rules

Often, reservoirs are operated in the short term to achieve specified storage targets. When current storage is below the specified target, releases are reduced. When current storage is above the specified target, releases are increased. The change in release is based on expected inflows and the difference between current storage and the storage target. These storage targets are not necessarily fixed, but can vary hourly, daily, or monthly and can depend on watershed and demand conditions.

Applicability and Optimality

Storage target rules are typically set to reflect a compromise in operations between conflicting objectives, such as water supply storage and flood control (Miquel and Roche, 1986). As such, particular storage targets and target rules are typically arrived at by simulation or optimization studies (USACE, 1977).

Application

The Columbia River System has an elaborate set of storage target rules used for operating individual reservoirs during short time periods within a refill season (USACE, 1985). An example of a seasonal storage target rule appears in Figure 2.12. Note that actual storage for that season imperfectly tracks the storage target.

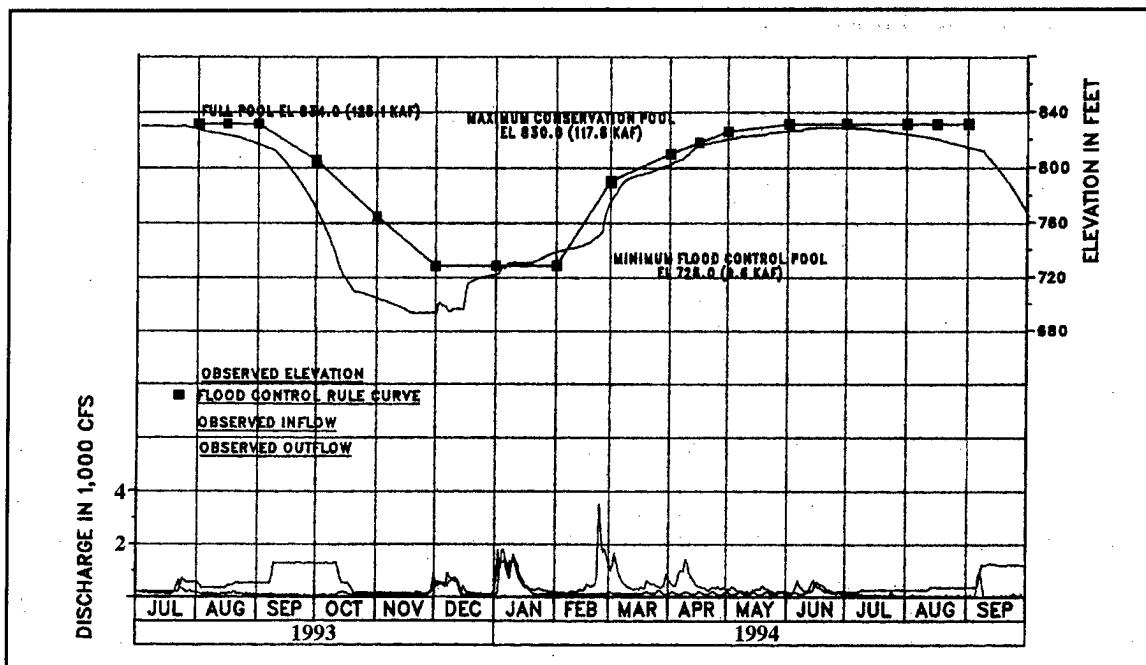


Figure 2.12: Storage Target Rule and Actual Operations

Storage target curves often vary with hydrologic conditions. For example, the absence of snowpack or dry watershed conditions can encourage raising storage targets, since under these conditions flood control is of less immediate concern. Similarly, high prices for hydropower can encourage reduction in near-term storage targets, allowing the system to take advantage of this opportunity for additional revenue. The Tennessee Valley Authority sets intermediate guide curves or targets, which vary with the price of energy (USACE, 1985).

Modeling and Implementation

Simple storage target rules are implemented within a time-step similarly to fixed storage rules, using the following equation:

$$R_t = \text{Max}(\text{Min}(S_t - S_{Tt} + I_t, R_{maxt}), 0)$$

where

R_t target reservoir release,
 S_t current reservoir storage, time t,
 S_{Tt} storage target for time t,

I_t expected net reservoir inflow, and
 R_{maxt} maximum allowed release for time t.

Typically releases are accelerated when storage is above the target, and curtailed when storage is below the target. USACE (1985) discusses the implementation and use of storage targets for hydropower and other reservoir purposes in detail.

Comments

Storage target rules are typically adjusted seasonally and within seasons in response to local and current conditions. This often requires extensive seasonal operating rule studies by simulation and/or optimization techniques. However, sometimes long-term studies are used to produce sets of storage target rules applicable for different year classes. Storage target rules are easy to express and employ in real time studies and allow disaggregation of a multi-reservoir system into individually operated reservoirs.

2.2.8 Complex Release Rules

When reservoir operating policies depend on multiple parameters, release decisions can become quite complex. They may involve the month of the year, current storage, expected inflows at upstream and downstream locations, or other such factors. These differ from simple release rules which may incorporate many of these factors but ultimately only depend on one or two of them in actual implementation. Complex release rules are used when simpler release or storage rules do not adequately meet desired objectives.

Applicability and Optimality

Reservoir system characteristics may be such that simple operating rules do not adequately meet system hydropower, water supply, or other system demands. Complex release rules add flexibility to the operating policy by incorporating more elements into the release decision. For certain reservoir systems increased complexity results in a more nearly optimal operation.

The establishment of complex release rules is typically done using simulation and/or optimization studies. Based upon such studies, release rules of often complex form can be found which result in desirable operations (USACE, 1994b).

Application

As a result of an HEC-PRM analysis on the Missouri River System (USACE, 1994b; Lund and Ferreira, 1996), a complex rule form has been suggested for the optimal operation of Oahe reservoir, one of six main stem reservoirs. The resulting rule varies releases based on inflows at various locations above and below Oahe reservoir. The inflow locations and their levels involved in the release rule vary by month. The rule for January is presented below, with locations as shown in Figure 2.13.

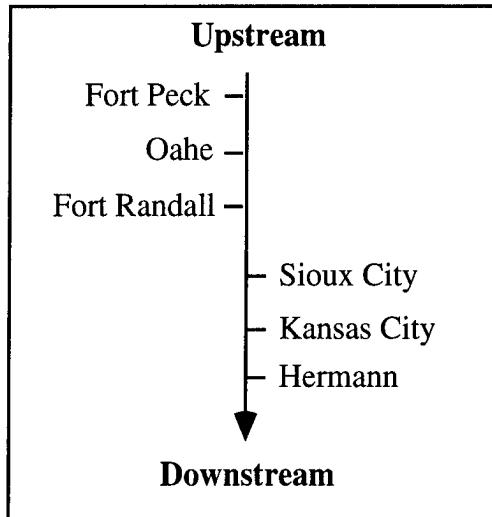


Figure 2.13: Inflow Locations Involved in an Oahe Reservoir Release Rule

```

If (FtP-OAH) Inf Dec < 550) and (60 < FtR-SUX) Inf Dec < 109)
or ((FtR-HER) Inf Dec < 870) and (100 < FtR-SUX) Inf Dec < 160)
    then Oahe Release = 553 - (FtR-SUX) inf KAF
else if ((FtR-HER) inf < 1900) and (800 < (FtR-KAN) inf < 1100)
    then Oahe Release = 1100 KAF
else if ((FtP-OAH) inf < 630) then Oahe Release = 553 - (FtR-SUX) inf KAF
else Oahe Release = 2000 KAF

```

where

(FtP-OAH) Inf	Combined January tributary inflow at all nodes between Fort Peck and Oahe
(FtP-OAH) Inf Dec	Combined tributary inflow at all nodes between Fort Peck and Oahe in previous December
(FtR-SUX) Inf	Combined January tributary inflow at all nodes between Fort Randall and Sioux City
(FtR-KAN) Inf Dec	Combined tributary inflow at all nodes between Fort Randall and Kansas City in previous December
(FtR-HER) Inf	Combined January tributary inflow at all nodes between Fort Randall and Hermann.

All values are in KAF per month.

Modeling and Implementation

The specific implementation of a complex operating rule depends on the particular characteristics and constraints of the reservoir system. It commonly consists of a series of if-then-else statements.

2.3 Rules for Reservoirs in Series

The next class of operating rules involves in the operation of reservoirs in series as shown in Figure 2.14. These rules determine the distribution of storages and releases necessary to meet the various demands imposed on the system. The rules discussed here are organized by various operating purposes.

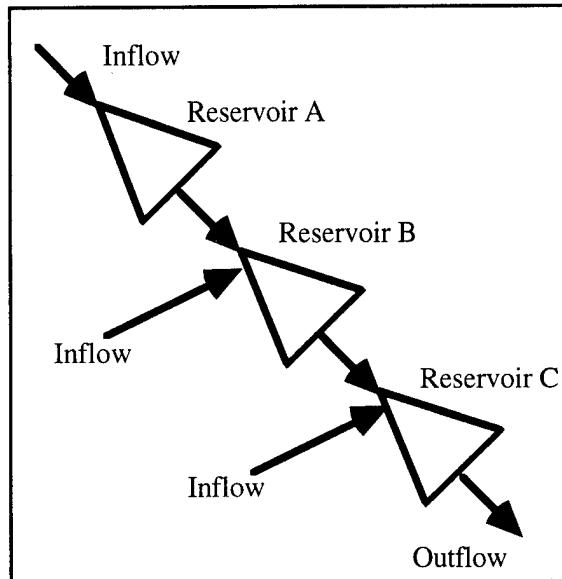


Figure 2.14: Reservoirs in Series

2.3.1 Water Storage Rules

For a system of reservoirs providing water supply a reasonable objective is to maximize the amount of water available, which is the same as minimizing spilled water. The resulting rule for single-purpose water supply reservoirs in series is simply, "Fill the higher reservoirs first, and the lowest last."

Applicability and Optimality

The likelihood and severity of shortages is reduced by preventing any water from leaving the system as uncontrolled and unproductive spills. For reservoirs in series, where there are intermediate inflows, the probability of spill from the system is minimized by first filling the uppermost reservoirs and retaining storage in the lower reservoirs to capture potentially large flows and reduce the likelihood of spills from the system. Spillage from any but the lowest reservoir is then captured by the next lowest reservoir.

Modeling and Implementation

For the case of a system containing one upstream reservoir j and one downstream reservoir i , the following pseudo-code fills the upper reservoir preferentially:

```

If ( $S_{jt} < K_j$ ) then
    if ( $S_{jt} - K_j + I_{jt} \leq 0$ ) then
         $R_{jt} = R_{jt,min}$ 
    else  $R_{jt} = S_{jt} - K_j + I_{jt}$ 
else  $R_{jt} = I_{jt}$ 

```

where

j	upstream reservoir,
S_{jt}	storage in reservoir j in period t ,
K_j	capacity of reservoir j ,
I_{jt}	uncontrolled net inflow into reservoir j in period t ,
R_{jt}	release from reservoir j in period t , and
$R_{jt,min}$	minimum allowable release from reservoir j in period t .

Alternatively, a linear-program type of rule can be formulated, as is done in the following section on flood control storage allocation for reservoirs in series.

This rule is particularly applicable to the refill season of the system, where inflows are in excess of downstream water supply demands. During the draw-down season, where system inflows are less than demands, the system should be drawn down in order of the upstream reservoirs first, to provide storage to accommodate potentially excess intermediate inflows or an early onset of the refill season.

Comments

For a series of reservoirs in series serving water supply as a sole purpose, the above rule seems universal. The sole exception to this might be where higher reservoirs suffer from higher rates of water loss due to evaporation and seepage (Kelley, 1986). In this case, any increased evaporation or seepage from higher reservoirs would have to be weighed against the increased potential for loss due to spill from concentrating storage at lower elevations.

2.3.2 Flood Control Rules

For reservoirs in series with intermediate inflows and storage serving solely for downstream flood control, it is optimal to regulate floods by filling the upper reservoirs first and emptying the lower reservoirs first.

Applicability and Optimality

Where reservoirs in series serve solely for flood control, the objective is to maintain as much control over flows entering the system above a critical flood-prone reach as possible. These reservoirs not only receive outflows from upstream reservoirs but also collect local runoff and tributary inflows. Flood control is exercised by maximizing available flood control storage at locations closest to the critical reach. Reservoir storage at locations closest upstream from the critical flood control reach always provides a higher level of flood controllability than any other reservoir (Mariën, et al., 1994). These are typically the lowest reservoirs in the series. Thus, for single-purpose flood control storage in a series of reservoirs, it is best to fill the higher reservoirs first and empty the lower reservoirs first.

An exception to this rule can be where the outflow capacity of the lower reservoir is restricted. Here, it can be better to fill the lower reservoir first to increase head on the outlet, thereby increasing release capacity from the entire system to the downstream channel capacity (USACE, 1976).

Application

An application of this rule can be found in a large Brazilian multi-reservoir system (Kelman, et al., 1989).

Modeling and Implementation

For the case of a system containing one upstream reservoir j and one downstream reservoir i , solving the following linear program simultaneously minimizes downstream flows above a given flow target Q_{Tt} , fills the upper reservoir preferentially, and empties the lower reservoir preferentially. The linear program would be solved as part of a simulation model for each time-step.

$$\text{Min } z = F_t + \epsilon S_{it} + \epsilon^2 S_{jt}$$

Subject to:

$$\begin{aligned} S_{j,t-1} + I_{jt} - Q_{jt} &\leq K_j \\ S_{i,t-1} + I_{it} + Q_{jt} - Q_{it} &\leq K_i \\ F_t &= Q_{it} - Q_{Tt} \\ &\{ \text{other reservoir operation constraints} \} \end{aligned}$$

where

F_t flood flows in excess of downstream flow target,
 ϵ an arbitrary constant, $0 < \epsilon < 1$,
 S_{it} storage in reservoir i at the end of the time step,
 $S_{i,t-1}$ initial storage in reservoir i ,
 K_i total flood control storage capacity of reservoir i ,
 I_{it} direct inflow into reservoir i ,
 Q_{it} release from reservoir i , and
 Q_{Tt} the target downstream flow.

In this linear program, Q_{it} and Q_{jt} are the major decision variables, which also define S_{it} , S_{jt} , and F_t . Storage capacities and initial storages are known. The constant ϵ is positive, but arbitrarily less than one. Inflows I_{it} and I_{jt} must be forecast for each time-step. The storage targets S_{it} and S_{jt} can be used as real time operating storage targets. An advantage of this linear program is that other reservoir operation constraints can be included in the constraint set. Alternatively, the solution of this linear program can be replaced with a large number of if-then-else statements to solve for each reservoir's release.

Comments

The operation of reservoirs in series for flood control is fairly complementary with water supply operations, at least in regard to the preferred location of storage. The maintenance of water supply storage still preys on the absolute flood control capability of a system, and vice versa. Nevertheless, where a given amount of flood control storage is to be allocated among a series of reservoirs with intermediate inflows, it is preferable to have that empty storage reserved, as much as possible, in the most downstream reservoirs above the most important flood control sites. USACE (1976) presents methods for allocating flood control space in reservoirs in series with flood control locations both downstream of the system and between reservoirs.

2.3.3 Hydropower Rules

Hydropower rules for reservoirs in series vary depending on whether the system is in the refill season or its drawdown season. During the refill season, the problem is attempting to maximize the storage of energy at the end of the season. During the drawdown season, the objective is to maximize hydropower production. Different rules are employed for each season. A difficult problem is the transition between seasons.

2.3.4 Energy Storage Rules

The objective of the energy storage rule for reservoirs in series is to maximize the total energy stored at the end of the refill season. It is a single-purpose rule. Here, the refill season is defined as the season when system inflows exceed those needed to meet water supply or hydropower production demands. The energy storage rule for reservoirs in series is to always fill the upper reservoirs first.

Applicability and Optimality

To maximize the energy stored for a future time, water storage typically should be shifted towards the upstream reservoirs. Water stored at higher elevations has a higher energy content (kilowatt-hours/acre-ft stored) than water stored at lower elevations. This is particularly true for water stored in reservoirs in series, where water eventually released from upper reservoirs generates hydropower at the lower reservoirs as well. Moreover, any overestimation of inflows into upper reservoirs, results in spills available for capture in space available in lower reservoirs. Kelman, et al. (1989) mathematically examine the allocation of energy storage and flood control storage capacity in complex multi-reservoir systems. Their results will often indicate the compatibility of the desirable distribution of energy and flood control storages in such two-purpose systems.

When it comes time to actually produce the necessary energy, hydropower rules to maximize power production may be employed. Downstream and smaller-capacity reservoirs are refilled. Releases from the upper reservoirs required to increase downstream storage are used by lower plants to generate power.

Application

The primary operating objectives of Tennessee Valley Authority's river control plan are flood control, navigation, and power generation (USACE, 1985). During the wet season, high levels of storage are maintained in tributary reservoirs at high elevations while the power load is carried by the main-river plants. This corresponds to the maximization of energy storage. Stored energy is used in the dry season when the tributary reservoirs are drawn heavily, providing released water to all downstream plants. Water from Fontana Reservoir, one of the tributary reservoirs at high elevation, is used by 13 downstream plants (Barrows, 1948).

Modeling and Implementation

Upstream reservoir storage is increased at the expense of downstream storage. For the case of a system containing one upstream reservoir j and one downstream reservoir i , implement the following pseudo-code:

```
If ( $S_{jt} < K_j$ ) then
    if ( $S_{jt} - K_j + I_{jt} - EV_{jt} \leq 0$ ) then
         $R_{jt} = R_{jt,min}$ 
    else  $R_{jt} = S_{jt} - K_j + I_{jt} - EV_{jt}$ 
else  $R_{jt} = I_{jt} - EV_{jt}$ 
```

where

j upstream reservoir,
 S_{jt} storage in reservoir j in period t ,
 K_j capacity of reservoir j ,
 I_{jt} uncontrolled inflow into reservoir j in period t ,
 EV_{it} evaporation from reservoir j in period t ,

R_{jt} release from reservoir j in period t , and
 $R_{jt,min}$ minimum allowable release from reservoir j in period t .

Comments

Fortunately, energy storage rules and water supply storage rules for reservoirs in series are quite compatible for the refill season, at least in terms of where storage is preferred in the system and their general intent to accumulate the maximum amount of water. However, with the coming of the drawdown season, hydropower production rules are required.

2.3.5 Hydropower Production Rule

The hydropower production rule attempts to maximize hydropower generation during a single time step, given a total storage target for the system. In essence, this involves the problem of how a given total storage should be allocated to maximize hydropower production. In general, the rule favors allocation of storage to those reservoirs which create a higher head per acre-ft of storage, have higher generation efficiencies, and have higher releases, since hydropower production is the product of head, efficiency, and release. A variant of this rule is the Corps' "storage effectiveness index" (USACE, 1985). This rule would typically apply to the drawdown season.

Applicability and Optimality

The maximum amount of power that could possibly be generated in a reservoir system would be generated when head levels in all reservoirs are at their highest. This is the approach used in single reservoir operation as well where storage levels are fixed at maximum levels. Where the amount of water stored in the system is limited, perhaps due to other operating purposes or variations in hydrology and demands, then the problem becomes one of allocating a limited amount of storage among the individual reservoirs to maximize hydropower production. This hydropower maximizing water storage allocation depends on reservoir capacities, inflows, efficiencies of energy production, and the total amount of water (or energy) to be stored.

Water is often stored first in smaller reservoirs, since head in relatively large reservoirs increases less per acre-ft of additional storage than in most small reservoirs. This artifact of geometry and topography is illustrated in Figure 2.15, where the volume of water needed to increase head by an amount x in the lower reservoir, V_2 , is much less than the volume needed to achieve an equivalent increase in head for the upper reservoir, V_1 .

Another consideration is release flow rate. All else being equal, hydropower production is maximized by allocating available stored water to reservoirs with the greatest release rates. Thus, for reservoirs in series with intermediate inflows, it is often desirable to maximize storage in the lower reservoirs first. Reducing storage in larger capacity reservoirs allows reservoirs with relatively little capacity to be operated at high storage levels.

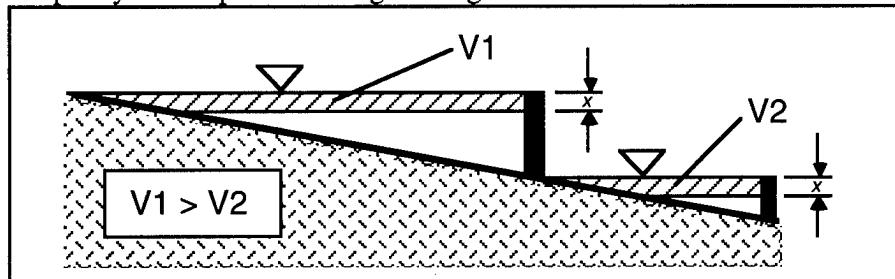


Figure 2.15: Change in Head With Varying Capacities

The relative location of reservoirs affects the amount of inflows that can potentially flow through reservoir turbines. Downstream reservoirs receive more direct and indirect inflows than do upstream reservoirs. Storage in downstream reservoirs is therefore kept at high levels to take advantage of the increased flows.

Lastly reservoirs with higher generation efficiencies should be maintained at higher levels of storage at the expense of reservoirs with lower efficiencies. The combination of reservoir capacity, amount of total inflows, and power generation efficiencies determines the overall caliber of the reservoir in producing power. When reductions in storage are necessary, they are made from those reservoirs with the least current ability to produce power. Conversely, if an increase in storage can be made, water should be stored in reservoirs with the greatest ability to produce power at the time the storage decision is made.

The interaction of these factors is examined mathematically in Appendix 2.1. The result is to calculate the following ratio for each reservoir i at each simulation time-step:

$$V_i = a_i e_i \left(\sum_{j=1}^i I_j \right),$$

where

- V_i increased power production per unit increase in storage
- a_i the unit change in hydropower head per unit change in storage (the slope of the head-storage curve),
- e_i the power generation efficiency of reservoir i , and
- I_j the direct inflows into reservoir j , for all reservoirs upstream of reservoir i .

Here reservoir 1 is the uppermost reservoir in the series of reservoirs. The hydropower production rule seeks to maximize storage in reservoirs with the highest values of V_i . Reservoirs are ordered in terms of their values of V_i , and are filled from highest to lowest value of V_i until the total water storage target is met.

This approach is conceptually similar to the Corps' "storage effectiveness index" (USACE, 1985). The calculation and use of the "storage effectiveness index" is presented in Appendix 2.3. For the drawdown season, these rules can be employed for systems of reservoirs in parallel, in series, as well as mixed systems.

Where the total storage constraint is desired to be in terms of energy storage, rather than water storage, then the more elaborate linear programming approach presented in Appendix 2.1 is required.

Application

Many systems of reservoirs in series maintain their lower, smaller reservoirs full for hydropower production. This is the case for the Missouri River System and the Columbia River System. Often, maintaining high storage levels in the lower reservoirs also aids navigation through the lower reaches of the reservoir system, as in the lower Columbia River System.

Modeling and Implementation

The hydropower production rule for reservoirs in series can be implemented directly using the values of V_i presented above, or by using the more elaborate linear programming formulations of the problem presented in Appendix 2.1. This rule should work best when employed for individual time-steps, where small changes in storage are anticipated.

Comments

Hydropower rules to maximize energy production are well suited for use during the drawdown season when energy demand is high in many parts of the country. During the refill season however hydropower rules that maximize energy storage may be more appropriate. The hydropower production rule and its "storage effectiveness rule" cousin, presented in Appendix 2.3, can also be applied to more complex reservoir systems during the drawdown season (USACE, 1985).

2.4 Rules for Reservoirs in Parallel

The operation of reservoirs in parallel, Figure 2.16, differs from reservoirs in series in that downstream reservoirs cannot be used to capture additional water from underestimated flows or benefit from the transfer of water stored upstream if flows are overestimated. Nevertheless, special rules have been developed for systems of parallel reservoirs. Generally, these are rules for balancing the storage among reservoirs in parallel and are referred to here as "space rules." The following sections present "space rules" for water supply, energy storage, water quality, and flood control purposes. These rules typically apply to the reservoir system's refill season. Computational studies suggest that these space rule forms tend to work rather well over a wide range of conditions, perhaps because the response surface is flat for such storage allocation decisions (Sand, 1984).

In much of the literature "the space rule" is a more specific rule for operation of parallel reservoirs (Bower, et al., 1966; Sand, 1984), referred to here as the "equal ratio space rule". The change in terminology used here results from a desire to generalize the idea of operating rules for parallel systems, since most of these rules are really variations on more general rules for parallel systems (Sand, 1984).

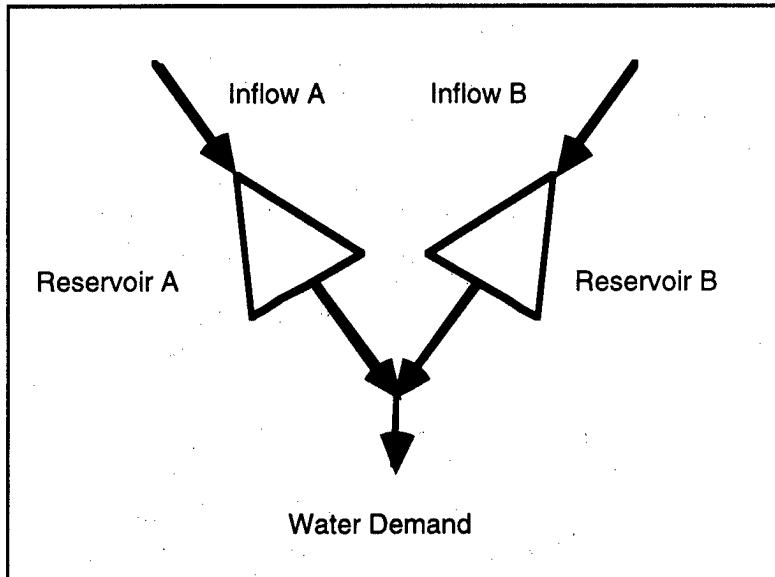


Figure 2.16: Reservoirs in Parallel

2.4.1 Water Supply, Energy Storage, and Water Quality Space Rules

Several space rules have been developed for water supply, energy storage, and water quality operations. For reservoir systems providing either water supply or hydropower production, a reasonable objective is to minimize expected shortages. The severity of shortages is reduced by avoiding any water leaving the system as uncontrolled and unproductive spills (Sand, 1984). Space rules prescribe ideal release or storage levels for reservoirs in parallel to avoid the inefficient condition of having some reservoirs full and spilling, while other reservoirs are unfilled (Bower et al., 1966). Several types of "space rules" have been developed for specific operating purposes for reservoirs arranged in parallel.

2.4.2 New York City Space Rules (NYC rules)

The NYC Space rules use the probability of spills rather than the direct amounts of physical spill in the minimization of expected shortages. When the probabilities of spilling at the end of the refill season are the same for every reservoir, it follows that physical spill also is minimized (See Appendix 2.2). Water supply shortfall is consequently minimized as well.

Applicability and Optimality

There are three important requirements for rigorous application of the NYC rule (Sand, 1984):

1. The system contains reservoirs operating in parallel;
2. The system provides for a single demand downstream of all reservoirs;
3. Expected shortages are to be avoided or minimized.

A modified form of the NYC rule also can handle situations where the unit value of water varies among different reservoirs but is constant in any individual reservoir. In terms of water supply, the quality of water might affect its unit value. For example, higher total suspended solids (TSS) concentrations correspond to greater treatment costs and lower desirability as a water supply source. For energy production the value of water is based on the maximum head of each reservoir; so water contained in reservoirs with greater head has proportionally greater value.

The application of the NYC rule depends on predicted inflows. It would thus seem that greater accuracy in these predictions would more likely yield optimal results. Since releases are recalculated at each period (usually monthly), a high degree of accuracy in predicted inflows is not critical in the early periods of the refill season though, reliable flow forecasts become more important as the chances of spill increase (Bower et al., 1966). Therefore, it is important to have enough historical and watershed data for probabilistic streamflow forecasts.

Optimality also depends on the coefficient of variation of mean monthly flows and the correlation between flows on adjacent streams (Bower et al., 1966; Sand, 1984). The NYC Space rule has been found to behave optimally or near-optimally for a wide variety of operating conditions and system configurations (Sand, 1984).

Application

The NYC space rule was first stated by Clark (1950) regarding its application to the New York City water supply system, "In operating this system an attempt is made to have the storage in each of the watersheds, at all times, fall on the same percentage year." The draw from each reservoir is adjusted to equalize the probability of refill by the end of the refill season, about June 1 (Clark, 1956).

Modeling and Implementation

The general form of the NYC rule equates the probabilities of spill at the end of the refill season adjusted by the unit value of water for each reservoir:

(1) $h_i \Pr[CQ_i \geq K_i - S_{fi}] = \lambda$ for all i where

- h_i unit value of water in reservoir i ,
- CQ_i cumulative inflow to reservoir i from the end of the current period to the end of the refill season,
- K_i storage capacity of reservoir i , assumed to be the same in every period, and
- S_{fi} end-of-period storage for the current period for reservoir i .

The values of h_i depend on water quality and energy storage issues as described in the following subsections. Historical data should be examined to determine the cumulative inflows, CQ_i . Release for the current period, R_i , is determined by taking the initial storage, S_{oi} , plus expected inflow for the current period, $E[Q_i]$, and subtracting the end-of-period storage, S_{fi} , that satisfies equation (1):

$$R_i = S_{oi} + E[Q_i] - S_{fi}$$

Water Supply

When the unit value of water, h_i , is the same among reservoirs providing water supply, h_i , is incorporated into the constant λ and thus drops out of the equation:

$$\Pr[CQ_i \geq K_i - S_{fi}] = \lambda \text{ for all } i$$

Water Quality

When the quality of water varies between reservoirs, such as varying TSS concentrations the probabilities of spill are adjusted by h_i , the marginal value of water use minus its treatment cost for each reservoir. Thus, if the marginal value of treated water use downstream is \$500/ac-ft and the cost of treatment for water from Reservoir 1 is \$50/ac-ft, $h_1 = \$445$.

$$h_i \Pr[CQ_i \geq K_i - S_{fi}] = \lambda \text{ for all } i$$

Energy Storage

For energy storage applications, the probabilities of the potential energy of spill are equated. Therefore the probabilities of spill are adjusted by h_i , the full head level of each reservoir.

$$h_i \Pr[CQ_i \geq K_i - S_{fi}] = \lambda \text{ for all } i$$

Comments

The NYC space rules apply to the refill season of systems of parallel reservoirs and attempt to minimize the expected value of spilled water. The primary difficulties are specification of inflow probabilities, computational use of the rule (but this should not be a major problem), and potentially in the absence of considering future refill season demands on the inflows into the system. This second problem may merit further exploration.

2.4.3 Equal Ratio Space Rule

Expected shortages may be reduced by minimizing volumes of spills. The equal ratio space rule takes this approach, although it is actually a special case of the more general NYC space rule (See Appendix 2.2). The equal ratio space rule seeks to leave more space in reservoirs where greater inflows are expected, or where greater potential energy of inflows are expected in the case of energy supply (Bower et al., 1966).

Applicability and Optimality

Since the equal ratio space rule is a special case of the NYC space rule, the same conditions for applicability and optimality apply. The NYC space rule becomes the equal ratio space rule when the distributional forms of inflows into each parallel reservoir are the same, with distributions scaled by their expected value, i.e., the distribution $f_i(CQ_i/EV(CQ_i))$ is identical for all reservoirs (Sand, 1984). This derivation appears in Appendix 2.2. The advantage of the equal ratio space rule over the NYC space rule is its slightly simpler computation. Like the NYC space rule, the spill minimizing objective implies that this rule is applicable to the refill season of the reservoir system. Like the NYC space rule, the equal ratio space rule has been found to behave optimally or near-optimally for a wide variety of operating conditions and system configurations (Sand, 1984).

Application

Johnson et al. (1991) examined the application of equal ratio space rules for operating the Central Valley Project in Northern California. In this system, power output is maximized while maintaining high levels of water supply reliability. This application appeared to offer improvements over simulated operation of the system.

Modeling and Implementation

The particular form of the equal ratio space rule depends on the reservoir purpose being examined for the system. Equal ratio space rules have been developed for both water supply storage and energy storage purposes.

Water Supply

For water supply purposes, implementing the equal ratio space rule consists of setting target storages in each reservoir so that the ratio of space remaining at the end of the current period to the expected value of remaining refill season inflow for each reservoir is identical (Johnson et al., 1991). This is expressed mathematically as,

$$\frac{K_i - S_{fi}}{EV(CQ_i)} = \frac{\sum_{i=1}^n K_i - V}{\sum_{i=1}^n EV(CQ_i)}, \forall i,$$

where V = the total water storage of the system at the end of the current time-step and all other terms are as defined in the NYC space rule.

Using the above equation, releases in a parallel system of reservoirs containing equally valued units of water are determined by the following procedure:

- 1) Determine the initial available storage, A_o , given capacity, K , and initial storage, S_o , for the reservoir, A_{oi} , and entire system, A_{oT} .

$$A_o = K - S_o$$

2) Determine available end of period storage for entire system, A_{fT} , using initial available storage, A_o , expected inflows for the period, Q_T , and target release, R_T .

$$A_{fT} = A_o - Q_T + R_T$$

3) To get available end of period storage for a particular reservoir, the space rule is applied. The ratio of available storage in each reservoir, A_{fi} , to available storage in the entire system, A_{fT} , equals the ratio of expected inflows into each reservoir, $EV(CQ_i)$, to expected inflows for the entire system, $EV(CQ_T)$, for the rest of the refill system.

$$A_{fi} = EV(CQ_i)/EV(CQ_T) A_{fT}$$

4) Determine the release for each reservoir for the period, R_i , using initial available storage, A_{oi} , final available storage, A_{fi} , and expected inflow for that period, Q_i .

$$R_i = A_{fi} - A_{oi} + Q_i$$

Energy Storage

When preventing energy spills, the equal ratio space rule equation used for water supply is modified by replacing reservoir capacities, available storage, and expected inflows with their potential energy counterparts. These consist of maximum potential energy that can be stored or the capacity of the turbines, available potential energy storage, and potential energy of expected inflows (Johnson et al., 1991). The substitution of these elements yields the following equation:

$$\frac{KE_i - E_{fi}}{EV(CE_i)} = \frac{\sum_{i=1}^n KE_i - E}{\sum_{i=1}^n EV(CE_i)}, \forall i,$$

where

- KE_i the maximum energy content of reservoir i ,
- E_i the target energy content of reservoir i at the end of the current time step,
- CE_i the cumulative energy inflow to reservoir i ,
- E the total target energy content of the reservoir system at the end of the current time step, and
- $EV()$ the expected value operator.

If releases or storages fall outside their permitted upper or lower bounds, the decision variables can be set to those bounds while the remaining variables are balanced according to the space rule (Stedinger et al., 1983).

Comments

Equal ratio space rules are simpler to implement than NYC space rules. However, they rest upon distributional assumptions which might not always hold. But the importance of these additional distributional assumptions can be tested for particular situations using long-term simulation modeling.

2.4.4 Linear Program (LP) Space Rule

The previous two forms of the space rule rely on historical data to determine expected inflows. The linear program (LP) space rule also uses past streamflow data, but rather than producing cumulative inflow distributions for CI_i , the data is entered directly into a linear program. Spill, or the value of spill, is minimized by considering all individual cumulative inflows from each period to the end of the refill cycle in past years. Compared to the NYC space rules, the primary advantage of LP space rules are their ability to incorporate other (linear) short-term reservoir operation constraints into the rule. Such additional constraints might include minimum or maximum flows downstream of each reservoir or required diversions below a subset of reservoirs.

Applicability and Optimality

The same conditions required for the NYC rule apply to the linear program space rule. Expected shortages are to be minimized and the system must contain reservoirs in parallel that serve the same purpose. The minimization of spill will depend on the representativeness of historical streamflow records. LP space rules are slightly more general than NYC space rules in that they can also incorporate other linear operating constraints in the setting of short-term storage targets for each reservoir. However, implementation of LP space rules require greater computational effort.

Application

Application of the LP space rule would require seasonal estimation of representative realizations of refill season cumulative inflows for each reservoir. These estimates would be required at each time-step. Such estimates could be taken directly from the historical streamflow records, be hydrologic modifications to historical streamflows, or be derived by use of synthetic hydrology. The records for each stream must be of the same length. For each time-step, the linear program below would be solved.

Modeling and Implementation

The LP space rule is implemented in the following form:

$$\min z = \sum_{j=1}^m \sum_{i=1}^n h_i L_{ij}$$

Subject to:

$$\sum_{i=1}^n S_{fi} = V$$
$$L_{ij} - E_{ij} = CQ_{ij} + K_i - S_{fi} \text{ for all } i \text{ and } j,$$

$$V = \sum_{i=1}^n (S_{fi} + Q_i) - D$$

plus any other linear constraints on present-period operations

where

m number of equally probable refill seasons;
 CQ_{ij} in hydrologic year j , the expected cumulative inflow to reservoir i from the end of the current period to the end of the refill cycle;

L_{ij} spill from reservoir i under hydrologic year j ;

E_{ij} empty storage capacity in reservoir i under hydrologic year j .

The values of h_i depend on whether water supply, water quality, or energy storage issues are being considered. Thus, water supply, water quality, and energy storage versions of the LP space rule can be developed.

2.4.5 Flood Control Space Rule

The approach taken for flood control in parallel reservoirs is to maintain a balance between reservoirs in terms of occupied capacities and flood runoff from drainage areas. If a reduction in outflows is required, it is made from the reservoir with the least percentage occupancy or smallest flood runoff. When an increase in releases is possible, it is made from the reservoir with the greatest capacity occupied or where relatively higher flood runoff is occurring. Higher releases

from reservoirs receiving greater flood-runoff may thus be counterbalanced by reducing releases from reservoirs receiving lesser runoff (Ghosh, 1986).

Applicability and Optimality

The intent of the flood control space rule is to operate the parallel reservoirs to balance the amount of flood control storage available, while maximizing undamaging releases from the system. While the principle of balancing flood control storage on parallel reservoirs should be clear, operation to meet this objective is not exact. If the objective were to minimize the expected value of damaging spills above the downstream channel capacity, then flood control space rules could be developed analogous to the New York City space rules. Unfortunately, the objective of flood control is more likely to be minimization of peak downstream flood flows during the refill season, where peak inflows to the system can arrive during a very short time. This situation is less rigorously represented by the NYC space rule approach.

Modeling and Implementation

A flood control space rule variant of the NYC space rule would set final storages in each period, S_{fi} , such that:

$$Pr[CQ_i \geq K_i - S_{fi}] = \lambda \text{ for all } i ,$$

$$\sum_{i=1}^n S_{fi} = V, \text{ and}$$

$$V = \sum_{i=1}^n (S_{fi} + Q_i) - Q_f,$$

where Q_f is the combined flood conveyance capacity of the reaches downstream.

USACE (1976) suggests the following method for allocating flood control space between two parallel reservoirs (A and B) with a single downstream flood damage location.

1) Route the project design flood and other observed floods with a maximum amount of runoff occurring above reservoir A and with maximum non-damaging releases from reservoir A. Allow reservoir B to make the remaining releases, up to the maximum non-damaging level. Plot the space required at reservoir A versus total space required.

2) Perform the same exercise for reservoir B, with the maximum design and observed flood flows entering above reservoir B. Plot the space required at reservoir B versus total space required.

3) The ratio for balancing flood storage between the two reservoirs should lie between these two curves.

Comments

While the concept of a flood control space rule appears conceptually sound, exact formulations remain unclear.

2.5 Miscellaneous Multi-Reservoir Rules

A general set of operating rules can be applied to systems containing single reservoirs, reservoirs in series, parallel, or a mixture of both.

2.5.1 Formal Optimization of Operations within a Time-Step

Many of the reservoir operating rules developed earlier in this chapter are derived from principles of optimization, and some require the use of some elementary optimization methods to assess desirable target storages or releases. However, optimization can be used explicitly in simulation modeling to optimize the operation of conveyance facilities and reservoirs within a time-step, and simultaneously allocate water among various water resource purposes to maximize net benefits or minimize net costs. This approach of "optimization within a time-step" avoids the assumption of perfect foresight present in long-term optimization, yet provides a realistic, and in principle optimal, operation of the system within each time-step. Between time-steps, simulation rules or optimization penalties on missing storage targets are used to maintain water in the system at appropriate levels.

Applicability and Optimality

Optimization can be employed within a time-step with relative ease. Typically, linear or network flow program solution algorithms are used as a substitute for traditional operating rules within each time-step. Penalties are assessed based on either maximizing economic performance or, more empirically, based on the relative desirability of meeting various target storages, flows, and releases.

Application

The "optimization within a time-step" approach has been employed widely (Jensen and Lund, 1993). WRMI (1994) applied this approach to water management for Alameda County, California. The approach has long been employed in the operation of the California Aqueduct (Chung, et al., 1989) and the modeling of urban water supply in Australia (Kuszera and Diment, 1988). The new system simulation model for the State of California (DWRSIM) also will use this general operating rule approach.

Modeling and Implementation

The use of optimization within a time-step requires the solution of a mathematical program for each time-step during the simulation period. The mathematical program consists of an objective function and constraint equations. Usually, these mathematical programs are adapted to take the following linear form:

$$\text{Minimize } z = \sum_{i=1}^n c_{qi} Q_i + \sum_{j=1}^m c_{sj} S_j$$

Subject to:

$$S_j = S_{0j} + I_j - Q_j, \forall j \text{ (mass balance for each reservoir)}$$

$$S_j \leq K_j, \forall j \text{ (maximum storage limits)}$$

$$Q_i \geq q_{min,i}, \forall i \text{ (minimum flows and releases)}$$

$$Q_i \geq 0, \forall i \text{ (non-negative flows and releases)}$$

$$S_j \geq 0, \forall j \text{ (non-negative storages)}$$

where

c_{qi}	unit penalty for flow on reach i,
Q_i	flow on reach i (and release from reservoirs),
c_{sj}	unit penalty for storage on reservoir j,
S_j	end-of-period storage on reservoir j,
S_{0j}	initial storage for reservoir j,
I_j	direct inflow to reservoir j,
K_j	storage capacity for reservoir j, and
$q_{min,i}$	minimum flow for reach i.

A very wide range of convex non-linear penalties can be incorporated into these formulations through the use of piece-wise linearization, permitting the setting of storage and flow targets, with increasing non-linear penalties for deviation. The basis for the penalties can be either economic (Jensen and Lund, 1993) or non-economic (WRMI, 1994; Chung, et al., 1989; Kuszera and Diment, 1988).

The use of optimization within a time-step usually is far more computationally efficient than long-term system optimization. Since the computation time for most optimization algorithms increases more than linearly with the number of decision variables and constraints, the time required to solve many small optimization problems (one for each time-step) is typically far less than the computation time required for solving one large optimization problem covering the entire period of analysis.

Comments

The use of formal optimization to operate the system within each time-step leads to a hybrid simulation-optimization model. This approach offers considerable flexibility in operation, since the model pursues the optimal solution (as it sees it) without regard to traditional operating rules. This use of optimization avoids much of the perfect foresight problems associated with long-term deterministic optimization. The difficulty arises mostly in the setting of target storages and penalties for reservoirs in the system, maintaining sufficient storage for droughts and/or sufficient storage capacity for future floods. As with traditional operating rules, the storage targets used in these models are usually determined experimentally, by varying target storages until satisfactory results are obtained.

This approach of optimization within a time-step also can be applied to single reservoir systems, particularly those which allocate water among multiple uses and users.

2.5.2 Storage Allocation Rules

Storage allocation rules are linear or non-linear "balancing" rules for distributing total storage targets among individual reservoirs in a multi-reservoir system. Given a total storage target or value for the entire system or a subsystem of reservoirs, the allocation rule determines storage targets for individual reservoirs (USACE, 1976; Loucks, et al., 1995; USACE, 1992). Storage allocation rules take the form of a storage allocation plot (Figure 2.17) (USACE, 1992) or a table (Table 2.3), where they are often known as "index levels" (USACE, 1976).

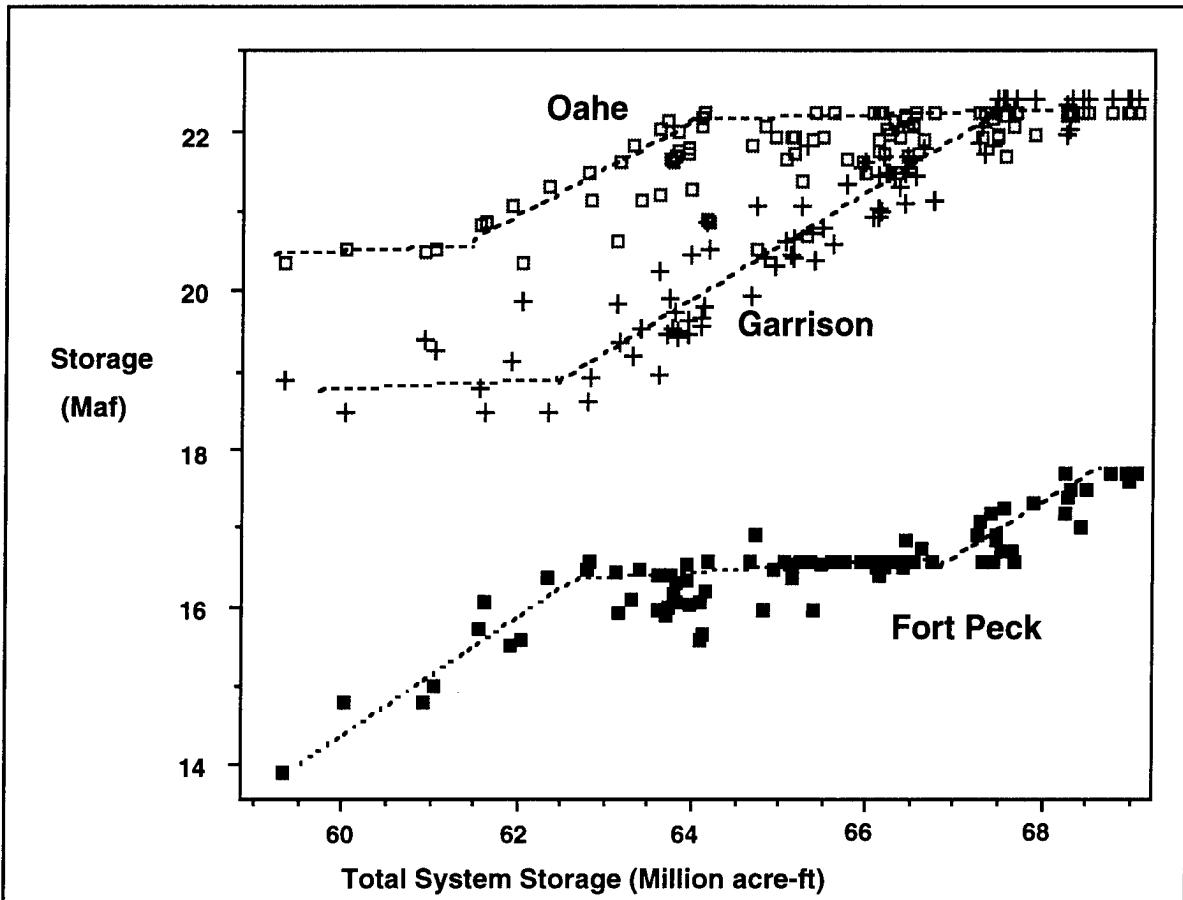


Figure 2.17: Storage Allocation Plot for Three Reservoirs in the Missouri River System for July.

Applicability and Optimality

Storage allocation rules can apply to any multi-reservoir system where a total system or sub-system storage target is to be distributed among individual reservoirs. Such balancing rules can be the expression of operating preferences supported by operating experience and simulation studies (Loucks, et al., 1995).

Storage allocation rules also can be inferred from reservoir optimization studies. Storage allocation rules, in the form of storage allocation plots such as Figure 2.17, were commonplace products from examination of HEC-PRM results for the Missouri River and Columbia River systems (USACE, 1992, 1994b, 1995, 1996). Storage allocation rules also can be developed based on detailed simulation studies (USACE, 1976).

Applications

Storage allocation rules clearly emerged from the analysis of HEC-PRM reservoir optimizations for the Missouri River and Columbia River systems (USACE, 1992, 1994b, 1995, 1996). In the case of the Missouri River system analysis, optimization results for 92-years of monthly operations showed patterns similar to that in Figure 2.17 for each month of the year. Later simulation studies showed that operation using monthly storage allocation rules closely matched HEC-PRM "optimized" storage operations for the largest of the systems six reservoirs (USACE, 1994b).

The results suggest a consistent non-linear or piece-wise linear allocation of total system storage among the upper three reservoirs of the Missouri River system. Referring to Figure 2.17, as the system is drawn down from a full condition, Fort Peck is drawn down first, until total system storage (including three downstream reservoirs) reaches roughly 67 MAF. Further drawdown is then taken primarily from Garrison, then Oahe, and then Fort Peck again. Refill storage allocation follows these rules in reverse.

Another expression of storage allocation rules in multi-reservoir simulation is the use of "index levels" (USACE, 1976; Loucks, et al., 1995), a tabular expression of potentially non-linear storage allocation plots. Storage allocation plots and rules often result as an expression of other rules, such as space rules (Sand, 1984).

Modeling and Implementation

Implementation of these rules for the Missouri River system was accomplished with a series of if-then statements (USACE, 1994b). For July, the pattern in Figure 2.17 was reduced to the following mathematical rules, where TS is total storage for the three reservoirs and units are in millions of acre-ft (MAF):

Fort Peck: If $TS < 56.5$ then Fort Peck Storage = $0.714 * TS - 23.9$
else if $TS < 60.5$ then Fort Peck Storage = 16.5
else Fort Peck Storage = $0.60 * TS - 19.8$

Garrison: If $TS < 56.2$ then Garrison Storage = 18.5
else if $TS < 61.3$ then Garrison Storage = $0.765 * TS - 24.5$
else Garrison Storage = 22.4

Oahe: $TS - (\text{Fort Peck Storage} + \text{Garrison Storage})$.

Such storage allocation rules were found for each month. These storage targets were used to determine releases from individual reservoirs.

Storage allocation rules also can be implemented through the use of "index levels" for the system. As explained in USACE (1976), "Index levels are integer numbers assigned to certain elevations in a reservoir. The levels are assigned in such a way as to control the 'balance' of the reservoir system. A system is 'in balance' when all reservoirs are at the same index level. ... In balancing levels among reservoirs, priority for releases is governed by the criteria that reservoirs at the highest levels at a given point in time are given first priority for making releases." An index level interpretation of the storage allocation plot in Figure 2.17 appears in Table 2.3 below. When total storage levels lie between index levels, target storages for the individual reservoirs are found by interpolation. The "index level" approach to reservoir balancing is applied in several reservoir simulation models (USACE, 1976; Loucks, et al., 1995).

Table 2.3: "Index level" Operation Table for Storage Allocation Plot in Figure 2.15, TS = total storage in the 3 reservoirs only

Index Level	Fort Peck	Garrison	Oahe	3-Reservoir Storage
6	17.7	22.4	22.2	62.3
5	16.7	22.4	22.2	61.3
4	16.5	21.8	22.2	60.5
2	16.2	18.5	21.5	56.2
1	0.714 * TS - 23.9	18.5	5.4+0.286 TS	TS < 56.2

Comments

Storage allocation rules in the form of "index levels" or storage allocation plots are simple ways of representing operating preferences for multi-reservoir systems. They can be applied to general multi-reservoir systems, including systems or sub-systems of reservoirs in series (Missouri River) and in parallel (Columbia River). However, the basis for storage allocation rules is largely empirical, resting on the results of optimization studies or operating preferences supported by actual experience or simulation studies.

2.5.3 Complex Release Rules

The complex release rules discussed earlier for a single reservoir can be extended to multi-reservoir systems. Here, releases for a given reservoir could be a complex function of the current month, as well as particular inflows, storages, or water use demands at any location in the system. These differ from simpler release rules which may incorporate many of these factors but ultimately only depend on one or two of them in actual implementation.

Applicability and Optimality

As with the single reservoir case, complex release rules are used when simpler operations do not adequately meet desired objectives. These rules add flexibility to the operating policy by incorporating more elements into the release decision. For some reservoir systems, increased flexibility results in a more nearly optimal operation. The optimality of such complex rules typically must be justified based on simulation testing.

Application

The case of Oahe reservoir on the Missouri River system, discussed under complex release rules for a single reservoir, also reflects the role of complex release rules in multi-reservoir systems. These rules were triggered by inflows at several locations throughout the system, and varied by month.

Modeling and Implementation

The specific implementation of a complex operating rule depends on the particular characteristics and constraints of the reservoir system.

Comments

Complex release rules are the last refuge of unruly reservoir operation results (USACE, 1992, 1994b).

2.6 Decomposition of Multi-Reservoir Systems

Multi-reservoir systems often can be decomposed into simpler sub-systems, particularly when their operation is dominated by a single operating purpose, such as downstream flood control, downstream water supply, or hydropower (Turgeon, 1981).

Decomposition of a complex reservoir system into simpler sub-systems is illustrated by Figure 2.18. In this case, a complex system is decomposed into a system in series and a system in parallel. For each sub-system, the relevant specific rules might be applied. For instance, for energy refill of a hydropower system of this configuration, energy storage space rules for reservoirs would be applied to the parallel system, consisting of reservoir C and the combination of reservoirs A and B. Energy storage rules for reservoirs in series would then be applied to the allocation between reservoirs A and B of their combined storage, established using the energy space rule. There appears to have been relatively little formal study of the utility and problems with decomposition of reservoir operations.

Finding reservoir operating rules by de-composition does not necessarily result in the best set of operating rules for a system. However, decomposition may be a promising approach to developing a strategy for system operations, or a structure for operating rules that can be refined by simulation or optimization studies.

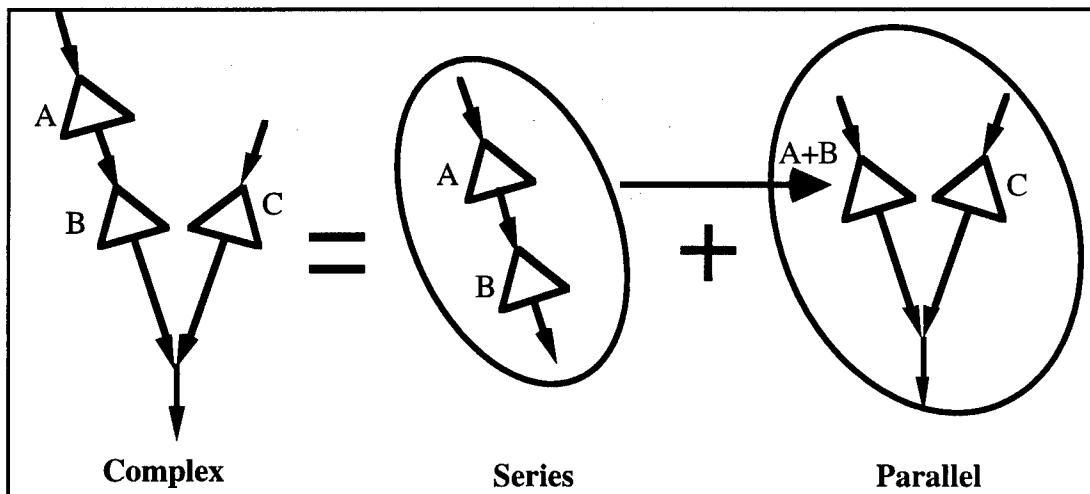


Figure 2.18: Decomposition of a Complex Multi-Reservoir System into a Sub-System in Series and a Sub-System in Parallel

2.7 Conclusions

Several conclusions are suggested from this review of reservoir operating rules.

1. A wide variety of reservoir operating rules are available for simulation modeling. These provide a great deal of flexibility in the specification of system operations under various flow, storage, and demand conditions.
2. A particular set of operating rules can be supported technically in a number of ways.
 - A) Some rules are based on simple engineering principles for reservoir operations, such as keeping reservoirs full for water supply or empty for flood control.
 - B) Several rules are derived from formal optimization principles, such as the New York City space rules and hydropower production and energy storage rules.
 - C) However, many rules are based largely on empirical or experimental successes, either from actual operational performance, performance in simulation studies, or optimization

results. These experimentally-supported rules are likely to be the most common for large multi-purpose projects.

3. Many opportunities exist for the use of formal optimization methods within reservoir simulation models. Examples include implementing storage allocation rules for reservoirs in series and in parallel, as well as general penalty-minimizing operations and allocating water among uses within a given time-step.

Appendix 2.1

Derivation of Hydropower Production Rules for Reservoirs in Series

Definition of Variables:

- P sum total of energy produced by all reservoirs,
- g unit weight of water,
- n number of reservoirs in system,
- H_i level of head in reservoir i ,
- S_i storage in reservoir i ,
- Q_i total inflow into reservoir i ,
- I_i direct inflows into reservoir i ,
- K_i storage capacity of reservoir i ,
- e_i efficiency of turbines in reservoir i ,
- a_i constant relating change in head with change of storage in reservoir i (for small changes in head),
- V_i change in overall power production, P with change in storage in reservoir i .
- γ unit weight of water

Note: reservoir $i = 1$ is most upstream reservoir in series.

Linear Programming Short Term Storage Allocation

The objective is to maximize hydropower production for one period, subject to inflow forecasts for each reservoir, reservoir storage capacities, and a total storage target. This is expressed mathematically below.

$$(1) \quad \text{Max } P = \gamma \sum_{i=1}^n H_i(S_i) Q_i e_i$$

Subject to:

$$(2) \quad Q_1 = I_1,$$

$$(3) \quad Q_i = Q_{i-1} + I_i, \quad \forall i > 1$$

$$(4) \quad S_i \leq K_i, \quad \forall i$$

$$(5) \quad \sum_{i=1}^n S_i = S$$

For short term allocation, the head-storage relationship can often be linearized, or $H_i(S_i) = a_i S_i$. The following linear program results:

$$(6) \quad \text{Max } P = \gamma \sum_{i=1}^n a_i e_i S_i Q_i$$

Subject to constraint Equations (2) - (5)

Where $H_i(S_i)$ is non-linear, $H_i(S_i)$ may be piece-wise linearized and solved with a linear program, since head-storage relationships for most reservoirs are concave.

Derivation of Hydropower Production Rule

The above linear programming formulation can typically be simplified, where the head-storage relationship can be linearized. Using objective function (6) and substituting in equations (2) and (3) results in the simpler linear program.

$$(7) \quad \text{Max } P = \gamma \sum_{i=1}^n a_i e_i \left(\sum_{j=1}^i I_j \right) S_i$$

Subject to:

$$(8) \quad S_i \leq K_i, \quad \forall i$$

$$(9) \quad \sum_{i=1}^n S_i = S$$

This problem is solved by finding the slope of the objective function with respect to storage for each reservoir,

$$(10) \quad \frac{\partial P}{\partial S_i} \alpha a_i e_i \left(\sum_{j=1}^i I_j \right) = V_i$$

Rule: Fill reservoirs in order of highest to lowest V_i until total storage, S , is filled.

Linear Programming Short Term Energy Storage Allocation

This problem is slightly modified from that above in that instead of seeking to maximize hydropower production subject to a given total water storage, it is desired to maximize hydropower production subject to a given total energy storage. For this problem, the objective and constraints in Equations 7 and 8 are re-stated below, and Equation 9 modified to an energy storage constraint.

$$(11) \quad \text{Max } P = \gamma \sum_{i=1}^n a_i e_i \left(\sum_{j=1}^i I_j \right) S_i$$

Subject to:

$$(12) \quad S_i \leq K_i, \quad \forall i$$

$$(13) \quad \sum_{i=1}^n E_i(S_i) = E,$$

where E is the total energy storage sought and $E_i(S_i)$ is the energy content of each reservoir as a function of its water storage.

For small changes in reservoir storage, the function $E_i(S_i)$ can probably be linearized into the form $E_i S_i$, allowing Equation 13 to be made linear and Equations 11-13 to be employed as a linear program to allocate storage among reservoirs in series to maximize hydropower production, subject to a total energy storage level.

Appendix 2.2

Derivation of Space Rules

These derivations of the New York City and Equal Ratio space rules are adapted from derivations presented by Sand (1984) and Johnson, et al. (1991). These derivations are further extended to examine energy storage and water quality applications.

Basic Derivation of the New York City Space Rule

Definition of Variables:

- z value of objective function,
- h_i unit value of water in reservoir i ,
- n number of reservoirs in the system,
- S_{fi} end-of-period storage for the current period for reservoir i ,
- S_{0i} beginning of current period storage for reservoir i ,
- CQ_i the cumulative inflow to reservoir i from the end of the current period to the end of the refill cycle,
- K_i storage capacity of reservoir i , assumed to be the same in every period,
- V total volume of water in storage at the end of the current period,
- I_i inflow to reservoir i for the current period,
- D demand for current period (release for current period), and
- $f_i(CQ_i)$ the probability density of CQ_i .

The objective of the New York City Space Rule is to minimize the expected value of total cumulative spill from all of the parallel reservoirs at the end of the refill season. This is reflected in the objective function in Equation 1. In Equation 1, the term h_i represents the relative value of water stored in each reservoir. Variation in the value of water can reflect variation in pumping costs, treatment costs, or energy content between the various reservoirs, as discussed later in the derivation. To implement this rule, this optimization problem is solved for each time-step during the refill reason.

$$(1) \quad \text{Min } z = EV \left(\sum_{i=1}^n h_i \min(0, S_{fi} + CQ_i - K_i) \right)$$

subject to the following constraints:

$$(2) \quad \sum_{i=1}^n S_{fi} = V$$

$$(3) \quad V = \sum_{i=1}^n (S_{0i} + I_i) - D$$

Constraint Equation 3 indicates that the total water available should equal the sum of available water (current storage plus current period inflows) minus downstream water demands for the current period.

Expanding the expected value function in the objective function (Equation 1) yields,

$$(4) \quad \text{Min } z = \sum_{i=1}^n h_i \left(\int_{K_i - S_{fi}}^{\infty} (S_{fi} + CQ_i - K_i) f_i(CQ_i) dCQ_i \right)$$

subject to constraint Equations 2 and 3.

The Lagrangian for this problem is:

$$(5) \quad L = \sum_{i=1}^n h_i \left(\int_{K_i - S_{fi}}^{\infty} (S_{fi} + CQ_i - K_i) f_i(CQ_i) dCQ_i \right) + \lambda \left(\sum_{i=1}^n S_{fi} - V \right)$$

$$= \sum_{i=1}^n h_i \left((S_{fi} - K_i) \int_{K_i - S_{fi}}^{\infty} f_i(CQ_i) dCQ_i + \int_{K_i - S_{fi}}^{\infty} CQ_i f_i(CQ_i) dCQ_i \right) + \lambda \left(\sum_{i=1}^n S_{fi} - V \right)$$

or

$$(6) \quad L = \sum_{i=1}^n h_i \left((S_{fi} - K_i) \left(1 - \int_0^{K_i - S_{fi}} f_i(CQ_i) dCQ_i \right) + \bar{CQ}_i - \int_0^{K_i - S_{fi}} CQ_i f_i(CQ_i) dCQ_i \right) + \lambda \left(\sum_{i=1}^n S_{fi} - V \right),$$

where

\bar{CQ}_i the expected value of cumulative inflows for reservoir i during the remainder of the refill season.

The first order conditions for solving this problem are:

$$(7) \quad \frac{\partial L}{\partial S_{fi}} = 0$$

$$= h_i \left(\left(1 - \int_0^{K_i - S_{fi}} f_i(CQ_i) dCQ_i \right) + (S_{fi} - K_i)(-f_i(CQ_i = K_i - S_{fi})) - (K_i - S_{fi}) f_i(CQ_i = K_i - S_{fi}) \right) + \lambda$$

or

$$(8) \quad h_i \left(1 - \int_0^{K_i - S_{fi}} f_i(CQ_i) dCQ_i \right) = \lambda,$$

or

$$(9) \quad h_i \Pr(CQ_i > K_i - S_{fi}) = \lambda, \forall i,$$

or

$$(10) \quad h_i \Pr(\text{any spill in reservoir } i) = \lambda.$$

This general result indicates that the storage targets for all reservoirs should have the same probability of spill, weighted by the value of water for each reservoir, h_i . The use of the New York City Space Rule for water supply, water quality, and energy storage purposes all follow Equations 9 and 10, with different interpretations of h_i .

New York City Water Supply Space Rule

For simple water supply purposes, h_i has the same value for all i , so Equation (9) becomes:

$$(11) \quad \Pr(CQ_i > K_i - S_{fi}) = \lambda, \forall i.$$

New York City Space Rule for Water Quality

For simple water supply purposes with important water quality differences (e.g., TSS) between reservoirs, h_i varies between reservoirs and can be interpreted as the marginal value of water use minus its water treatment cost for reservoir i . In this case Equations 9 and 10 remain the same, but with this net water value varying with spills of different water qualities.

New York City Space Rule for Energy Storage

Here, the objective is to minimize the expected value of potential energy spilled rather than physical water spilled. Here, h_i has the interpretation of the energy content of water stored in reservoir i . Equations 9 and 10 remain the same and applied with this interpretation.

Derivation of Equal Ratio Space Rules

Returning to Equation 9, the central result of the New York City Space Rule, the assumption is made that the distributions $f_i(CQ_i)$ have the same distributional form, except that they are scaled by the average cumulative flow of the basin, \overline{CQ}_i . Where this assumption holds, then the distributions,

$$(12) \quad f_i(CQ_i / \overline{CQ}_i) = f_j(CQ_j / \overline{CQ}_j),$$

for any two reservoirs i and j .

Where this is the case, the ratio $(K_i - S_{fi}) / \overline{CQ}_i$ becomes a standard deviate for all distributions, having the same probability of exceedence for all reservoirs. If this ratio is set so that it equals the same ratio at the basin-wide scale,

$$(13) \quad \frac{\sum_{i=1}^n K_i - V}{\sum_{i=1}^n EV(CQ_i)},$$

then the reservoirs are all balanced in terms of minimizing expected value of spill and maximizing capture of current inflows.

By replacing water inflows, water storage capacities, and water storage levels with energy inflows, energy storage capacities, and energy storage levels, the equal ratio space rule can be adapted to energy storage purposes much as the NYC rule can be adapted to other operating purposes (Johnson, et al., 1991).

Derivation of Linear Programming Space Rules

Derivation of the linear programming space rules begins with the Equations (1-3) used in the derivation of the NYC space rules. Additional constraints can also be added to the LP space rule problem, so long as the additional constraints are linear.

In this case the expected value operator in Equation 1 is replaced by use of the weighed summation of spill values that would result from each year of the historical record. Given historical streamflows of equal record length m for each of n reservoirs, m equally probable refill seasons can be inferred. This yields the following linear program:

$$\min z = \sum_{j=1}^m \sum_{i=1}^n h_i L_{ij}$$

Subject to:

$$\begin{aligned} \sum_{i=1}^n S_{fi} &= V \\ L_{ij} - E_{ij} &= CQ_{ij} + K_i - S_{fi} \end{aligned}$$

$$V = \sum_{i=1}^n (S_{fi} + Q_i) - D$$

plus any other linear constraints on present-period operations

where

- m number of equally probable refill seasons;
- CQ_j in hydrologic year j , the expected cumulative inflow to reservoir i from the end of the current period to the end of the refill cycle;
- L_{ij} spill from reservoir i under hydrologic year j ;
- E_{ij} empty storage capacity in reservoir i under hydrologic year j .

Appendix 2.3

"Storage Effectiveness Index" Rules for Hydropower Production

The "storage effectiveness index" has been developed by the U.S. Army Corps of Engineers for maximizing firm hydropower production during the drawdown season (USACE, 1985). For each reservoir, a "storage effectiveness index" is calculated for each time-step, using forecast inflows and power demands for the current time-step and remaining time-steps in the drawdown season. Reservoirs with a low index value are drawn down first.

Step 1: Find the firm energy requirement for the current time-step, E_f .

Step 2: Estimate the shortfall of firm hydropower production due to insufficient inflows to the system.

$$S_f = E_f - \frac{720}{11.81} \sum_{i=1}^n I_{Ui} H_i(S_i) e_i,$$

where,

S_f = energy shortage for the current time-step,

I_{Ui} = inflow upstream of reservoir i during the current time-step,

$H_i(S_i)$ = hydropower head as a function of reservoir storage for reservoir i ,

S_i = current reservoir storage for reservoir i ,

e_i = the hydropower production efficiency of reservoir i , and

the constant is a conversion factor for I_{Ui} in cfs, H_i in feet, and S_f in Kwh.

This assumes that all flow can be utilized through the turbines.

Step 3: For each reservoir, estimate the drawdown required for that reservoir to individually eliminate the shortfall.

$$S_f = \frac{720}{11.81(59.5)} \Delta S_i H_i e_i,$$

where,

$1/59.5$ = conversion of cfs to acre-ft draft per month,

ΔS_i = drawdown, in acre-ft., and

H_i = an average head corresponding to the drawdown (often found iteratively).

Solving for ΔS_i :

$$\Delta S_i = \frac{11.81(59.5)}{720} S_f / (H_i e_i).$$

Step 4: For each reservoir, estimate the energy loss in the remainder of the drawdown season due to a drawdown of ΔS_i during this time-step (month).

$$E_{Li} = \frac{720}{11.81} (C I_{Ui} + V_{pi}) H_i (S_i - \Delta S_i) e_i / 59.5,$$

where,

E_{Li} = drawdown season power loss due to drawdown of reservoir i by ΔS_i ,

CI_{U_i} = the cumulative natural inflow upstream of reservoir i for the remainder of the refill season, and

V_{pi} = the volume (acre-ft) of upstream storage to be emptied during the remainder of the drawdown season.

Step 5: Calculate the storage effectiveness ratio for each reservoir i :

$$SER_i = E_{Li} / S_f.$$

Reservoirs with the lowest ratios are to be drawn down first.

Chapter 3

Long-term Rules from Optimization Studies

3.1 Introduction

This chapter reviews approaches for developing long-term operating rules from deterministic optimization results. These approaches include those discussed in earlier HEC reports (USACE, 1992), as well as newer artificial neural network and statistical methods (such as principal component analysis).

There are five steps to developing operating rules from deterministic optimization results:

1. Formulation of the optimization model's objective function, capacity constraints, and mass balance constraints (USACE, 1991a, 1991b),
2. Selection of one or more appropriate hydrologies,
3. Solution of the deterministic optimization model (USACE, 1991a, 1993a, 1994a),
4. Inference of operating rules from the deterministic optimization results, and
5. Testing and refinement of inferred operating rules by simulation modeling (1994b).

This chapter focuses on steps 2, 4, and 5. Steps 1 and 3 are well covered elsewhere for applications of HEC-PRM (USACE, 1991a, 1991b, 1993a, 1994a).

3.2 Hydrologies for Rule Development

HEC-PRM and other optimization models prescribe a set of operating decisions for a given set of hydrologic inputs. There are a variety of approaches for establishing the hydrology to be used for such operating rule development. The selection of a hydrology or hydrologies should vary with the intent of the rule-making exercise.

Developing Rules for "Design Periods"

Sometimes, deterministic optimization is used to examine ideal reservoir operations during specific design conditions, essentially identifying ideal operation under a given hydrologic "design load". The most common "design events" are "typical", drought or "critical period," and "design flood" conditions. These periods are typically much shorter than the historical record and so require relatively little computation time. Operating rules for such special conditions are typically used qualitatively in a larger context to aid in the development of a more general set of operating rules.

"Typical" years for Operations Studies

A common approach for developing operating rules from deterministic optimization results is to specify a hydrology and water demands for a "typical" year or a set of typical years. Deterministic optimization is then used to find optimal operations for such years and these optimal results are then interpreted to find operating rules, often with the aid of simulation (King and Evenson, 1972). Rules developed by this approach may be informative, but will not be applicable to as wide a range of conditions as those developed by implicit stochastic optimization using a much longer streamflow record.

"Critical periods"

The performance of systems during simulations of critical drought periods is commonly used for establishing operating rules. The "critical period" for a reservoir system is typically defined as that period from the hydrologic record for which the system would have its minimum constant yield. This period can vary for different reservoirs and different sub-systems, and also can vary depending on whether the system of reservoirs is operated conjunctively or independently (USACE, 1985). In theory, it is also possible for different purposes of the system, such as energy yield, water supply yield, or water supply yields at different locations, to have different "critical periods." Thus, it is not uncommon for reservoir systems analysis to be conducted for more than one critical period. Traditionally, analyses of a system's critical periods is the basis for estimates of the firm energy and water yields of the system, used in contracting and operations. Clearly, both operators and users of reservoir systems have an interest in how the system is likely to perform under extreme circumstances, even though it is quite possible to imagine and expect worse streamflow events.

The use of critical periods as the hydrologic input for deterministic optimization models provides an idea of the best performance possible under the worst hydrology of record and often will provide insight into how the system should be operated for such extreme events. By bounding this extreme performance level, the optimization results provide operators and users with a surrogate for the system's "firm yield", even though such estimates are likely to be a bit optimistic due to the optimization model's perfect foresight. The operation policy suggested by the optimization model results should be taken as primarily advisory, given the potentially unique circumstances and perfect foresight assumed by the optimization model.

Where multiple critical periods are available and employed, and a relatively consistent operating policy is realized for all critically dry events, this would provide much stronger support for a potential drought operations policy. For this case, a major problem remaining would be when to declare the system in a state of drought.

"Design flood"

"Design floods" are often used to guide the specification of operating rules for flood conditions. "Design floods" can be defined based on historically realized floods or can be synthesized from climatic and rainfall-runoff patterns. For many small systems a 100-year flood hydrograph is often synthesized for reservoir operations studies. For most systems, one or more of the largest historical floods are also used to examine the flood control performance of potential operating rules. As with the use of "critical periods" in drought operations, the use of deterministic optimization for "design floods" provides an upper bound on the performance of the system under such conditions. In addition, if optimization results for multiple "design floods" of different characters show similar operating patterns, this constitutes evidence of desirable operating procedures for such flood events.

Rules Based on Long-Term Optimization

Long-term deterministic optimization studies generally are less intent on optimizing performance for specific short episodes of system operation (such as floods or droughts), but more concerned with improving the operation of the system on average over a representative range of wet, dry, and normal conditions. Such studies typically require much larger amounts of streamflow data. For these situations, the deterministic optimization model typically uses one or more sets of long streamflow record.

Historical period performance

The simplest approach to developing a hydrologic record for long-term deterministic operations is to use the historical streamflow record. Such records already exist for most large and medium-sized systems and often have been updated to reflect land-use changes and upstream depletions. This approach was employed successfully in application of HEC-PRM to develop rules for the Missouri River System (USACE, 1991a, 1994b), with 93 years of record, and Alamo Reservoir in Arizona (Kirby, 1994), with 103 years of record. Preliminary analyses for the Columbia River, with 50-years of record, also appeared to be adequate for examining strategic operating rules (USACE, 1995). Jettmar and Young (1975) found promising operating rules for a simpler system using relatively short (64-year) streamflow records. As time marches on, available streamflow records typically become longer and the use of the historical record becomes more attractive. In some cases, it might be desirable to extend the available streamflow record using longer precipitation or other climatic records.

Synthetic streamflow record performance

Synthetic streamflow records have been the traditional inputs for classical academic uses of implicit stochastic optimization (Young, 1966; Karamouz, et al., 1992) and some applied uses (Saad, et al., 1994). Statistical streamflow synthesis techniques can provide an arbitrarily large number of arbitrarily long synthetic streamflow records for use in deterministic optimization models (Linsley et al., 1982; Viessman, et al., 1977). These statistical synthesis techniques attempt to provide synthetic streamflow time series with the same statistical characteristics as relevant historical records. The primary advantages attributed to the use of synthetic records is their ability to produce floods and droughts likely to be worse than those present in shorter historical records, a more statistically representative range of streamflow events, and (where historical records are short) a streamflow of adequate length for system studies. Some disadvantages also can apply to the use of statistically synthetic streamflow records where they cannot be communicated acceptably well to engineering and decision-making users, are not methodologically clean or verifiable, or produce flood and drought events with uncommon qualities (Klemes, 1974; Eltahir, 1996).

Model Warm-Up and Cool-Down

"Warm-up" and "cool-down" periods are the beginning and ending periods of a model run which are particularly affected by the initial and final conditions of the model. An example of this phenomenon appeared in the original Missouri River System HEC-PRM model, where no final storage target or penalty was specified for the 93-year model run (USACE, 1991a, 1992). The result was that optimized releases from the system increased greatly during the last years of the model run, draining the system. There was no special penalty for leaving the reservoirs empty at the end of the run, so the model drained the system at the end to reduce water supply and hydropower penalties during the last few years. In this case, the last three years of the model run were eliminated from the analysis of results. The long length of this discarded record arose because of the unusually large storage capacity of the Missouri River System relative to its inflow (over 3 years of average inflow).

Similar problems can arise when the initial storages specified for the reservoir system model are far from the range of optimized results. It can take several time steps to raise or lower reservoir storages to desirable levels. Thus, for most systems, the first few months of optimization results might be unrepresentative for purposes of developing long-term operating rules. Since most reservoir optimization models begin with reservoir storages at fairly reasonable levels, the period of warm-up is usually fairly short.

The period of "cool-down" at the end of the streamflow record can be lessened or eliminated through several techniques. The simplest approach is to specify a representative end-of-period storage for each reservoir in the system. Problems with this approach are that the model may have to greatly disturb the near-end results to achieve this storage level, in cases where floods or droughts occur toward the end of the record and the specified storage might not actually be representative of optimized operations. In any case, this approach is likely to be superior to not specifying any end-of-period storages and has been successful (USACE, 1993a).

A more sophisticated approach is to place a penalty function on end-of-period storage (USACE, 1995). This approach requires that a target storage be set for the last period and that a penalty be set for failure to achieve this storage. These penalties often can be set reasonably to represent the value of stored energy (where hydropower is important) or stored water for the future draw down.

Summary

The selection of a set of hydrologic inputs for a deterministic optimization model is likely to vary with the intended use of the modeling study, the availability of representative historical records, and the assignment of beginning and ending conditions to the system. For application to particular conditions, such as floods, drought, or relatively near-term scheduling, relatively short periods of inflow are needed, but they must be representative of expected inflow events (often a problem for floods and droughts). For application to long-term strategic operating rule studies, a long, reliable historical record is probably best, but often will be unavailable.

3.3 Inferring Rules from Results

A variety of general approaches are available for discerning reservoir operation rules from deterministic optimization results. Variants of these approaches have been employed in previous optimization studies.

Each approach seeks to detect and substantiate a pattern in long-term optimal operations that can be reduced to "rules" triggered by the reservoir operator's current state of knowledge. Thus, operation rules must be based on known states such as: the current month, current storage, and current or forecast inflows. For the Missouri River system, some typical examples of operation rules would be:

- A storage rule based solely on the month,
"In February, keep Fort Randall storage at 3.5 MAF."
- A storage rule based on the current month and system storage,
"In July, if total storage > 64 MAF, keep 22 MAF in Oahe."
- A release rule based on system storage,
"In July, for total storage between 50.5 and 59.0 MAF maintain a flow of 25,000 cfs + 706*(Storage - 50.5 MAF) at Sioux City."

In each case, the rule determines the operating decision based on something known to the system operator.

The major difficulty in detecting these rule patterns in long-term optimization results is the amount of optimization result data available. For the case of the 90-year record used in the Missouri River exercise, a total of 12,960 optimal release and storage decisions were provided, in addition to input inflow data and data on consequent downstream flow consequences of release decisions. The approaches discussed below are those commonly employed to identify consistent trends in large amounts of data.

Intuitive Approaches

Intuitive approaches to discerning reservoir operation rules employ our innate and educated abilities as engineers to detect significant patterns in data. For example, we all feel able to "see" when plotted data seem to fit a linear trend.

The use of intuition for identifying and substantiating apparent "rules" in optimization results is greatly aided by graphical and statistical tools. Descriptive statistics, histograms, scattergrams (data plots), and other techniques all present data in a form conducive to our "seeing" trends. Spreadsheets and statistical and data analysis software packages can be very valuable in quickly providing a wide variety of such displays and descriptive statistics to the rule-maker.

An educated intuition was the major approach used in developing operating rules from application of HEC-PRM to the Missouri River System (USACE, 1992, 1994b). For this application, graphical and statistical displays were essential to identifying most of the underlying trends in the optimization model results. For example, trends in the operation of the three lowest reservoirs in the system were found to consistently vary by month through the use of quartile plots of reservoir storage by month. An example of such a plot appears in Figure 3.1. In this graph, the operation of Fort Randall can be seen to vary rather consistently and with relatively little variability within each month. This simple plot was used to identify operations for three reservoirs, resulting in operating rules consisting of monthly storage targets.

More elaborate multi-reservoir operations for the Missouri River System were identified through the use of storage allocation plots, such as that in Figure 3.2. Here, total system storage is plotted against individual reservoir storages for each month (or over longer periods). As explained in Chapter 2, such plots can indicate the relative order of drawdown and refill of reservoirs in multi-reservoir systems and the allocation of total system storage among several reservoirs.

The utility of intuition in rule-making is limited by the intuitive abilities of the rule-maker and the complexity of the rule-making task. There may always exist a better pattern that a rule-maker is unable to "see." Also, different rule-makers might "see" different patterns. Finally, the complexity and quantity of the data may be difficult to present in a form conducive to intuitive rule-making.

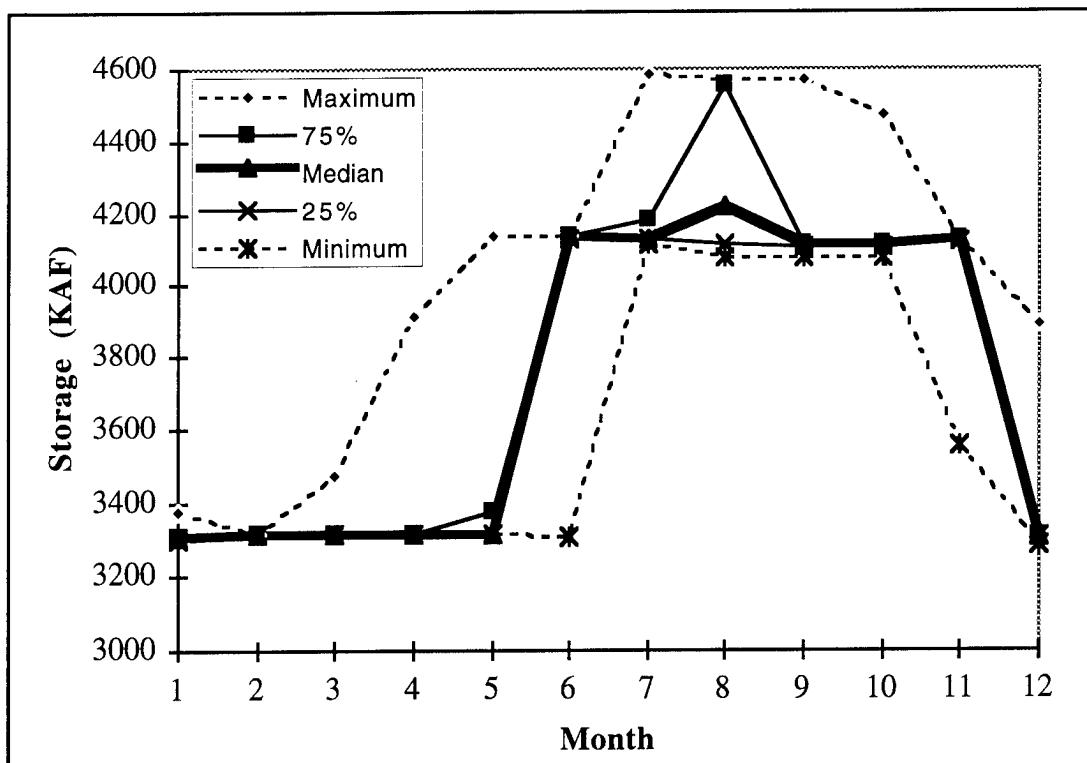


Figure 3.1: Quartiles of Seasonal Storage in Fort Randall under HEC-PRM Operations

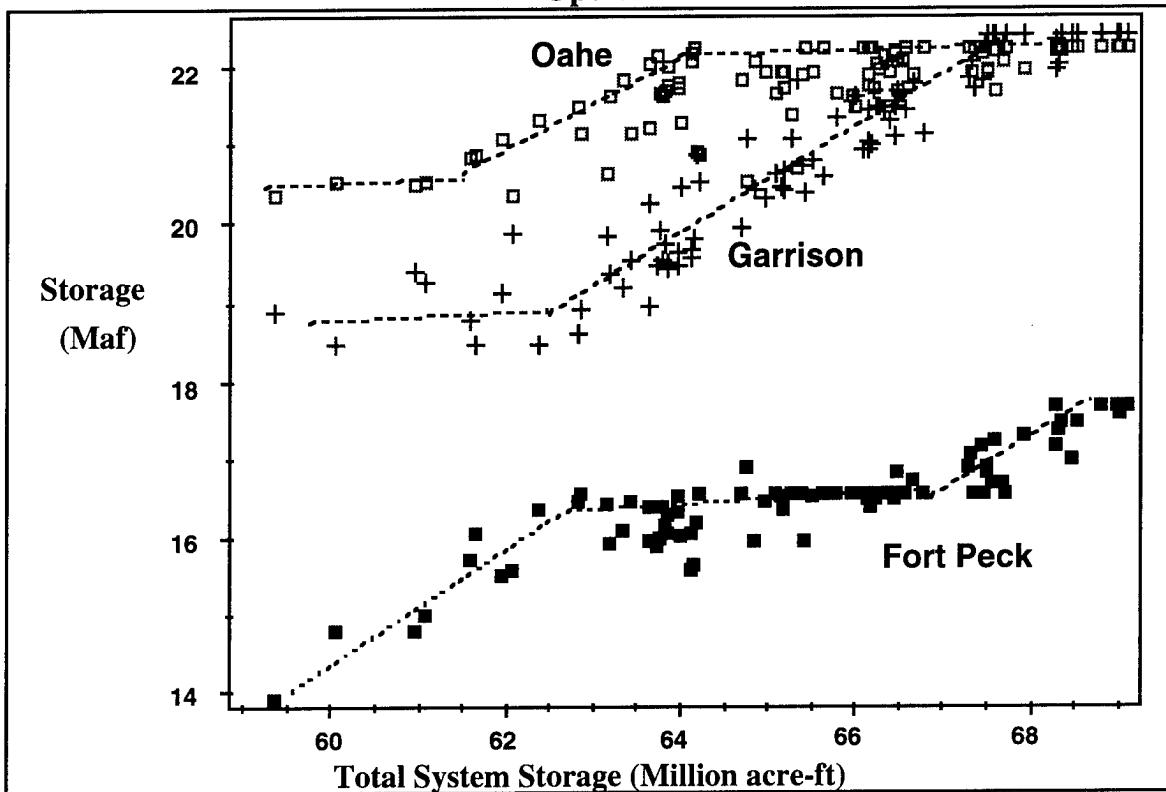


Figure 3.2: July Storage Allocation Results for the Missouri River System

Reservoir Operation Theory

Reservoir operation theory can be of great use in suggesting the forms of operating rules that might be fit to optimal operation results. Work on optimal rule forms and patterns can be particularly useful (Clark, 1956; Maass et al., 1966; Kelman, et al., 1989; Loucks and Salewicz, 1989; Johnson, et al., 1991). Some common examples of these optimal operating rule forms are:

- Space rules (Clark, 1956; Maass et al., 1966; Johnson, et al., 1991), which seeks to balance storage between reservoirs in parallel to minimize the likelihood of spills,
- Pack rules (Maass et al., 1966), which maintain storage at high levels as long as possible to increase hydropower heads and production, and
- Hedging Rules (Maass et al., 1966), which reduce reservoir releases early in a drought to reduce the risk of shorting more critical release uses later in a drought.

A review of reservoir operation rules and their theoretical background appears in Chapter 2.

Regression Approaches

Regression typically tries to develop equations which predict optimal decisions, such as releases, based on input data, such as current month, current storages, and forecast inflows. Regression techniques typically assume linear relationships between these variables and attempt to best "fit" the regression equation by finding parameters for the equations that satisfy some "fit" criterion, such as minimizing the sum of squared deviations between the optimal decisions and decisions predicted by the linear regression model.

Regression was first employed for developing reservoir operation rules from optimization results by Young (1966) and has been employed by others since (Jettmar and Young, 1975; Bhaskar and Whitlatch, 1980; Karamouz and Houck, 1982; Karamouz, et al., 1992). Before using regression to estimate an operating rule, specific dependent and independent variables must be defined. Independent variables would include those things known at the time of real operations, such as the current month, current storages, and current inflows. The dependent variable in the regression would be some operating decision which must be made, such as a release rate or a storage target. Given the relative ease of performing regression analysis with contemporary statistical packages, it is easy to explore a variety of dependent variables and several combinations of independent variables. The specification of independent and dependent variables is rather subjective, aided by intuition and judgment, reservoir operation theory, simulation results, and previous regression results.

Most use of regression for developing reservoir operation plans has been for single reservoirs (Young, 1966) or small multiple reservoir systems with a single operating purpose (Bhaskar and Whitlatch, 1980; Karamouz, et al., 1992). Usually, this approach has been used to develop an operating rule relationship between reservoir release and reservoir storage (or storage plus current inflow).

For larger reservoir systems, such as the Missouri River system, there are many possible sets of independent and dependent variables. The operation of multi-purpose reservoir systems, where the optimal operation is driven both by storage, release, and downstream flow values is also less likely to be revealed by simple linear relationships, as shown in the results appearing in Figure 3.3. Quentin Martin (unpublished memo, 1992) suggests the use of piece-wise linear regression for establishing particular operating rules from HEC-PRM results. In addition to the engineering judgment, intuition, theory and other aides to specifying independent variables, step-wise multiple regression can be useful in assessing which of many possible independent variables tend to best explain variation in a particular release rate or storage level. This secondary approach to regression was found to be quite useful for Missouri River system operation (USACE, 1992).

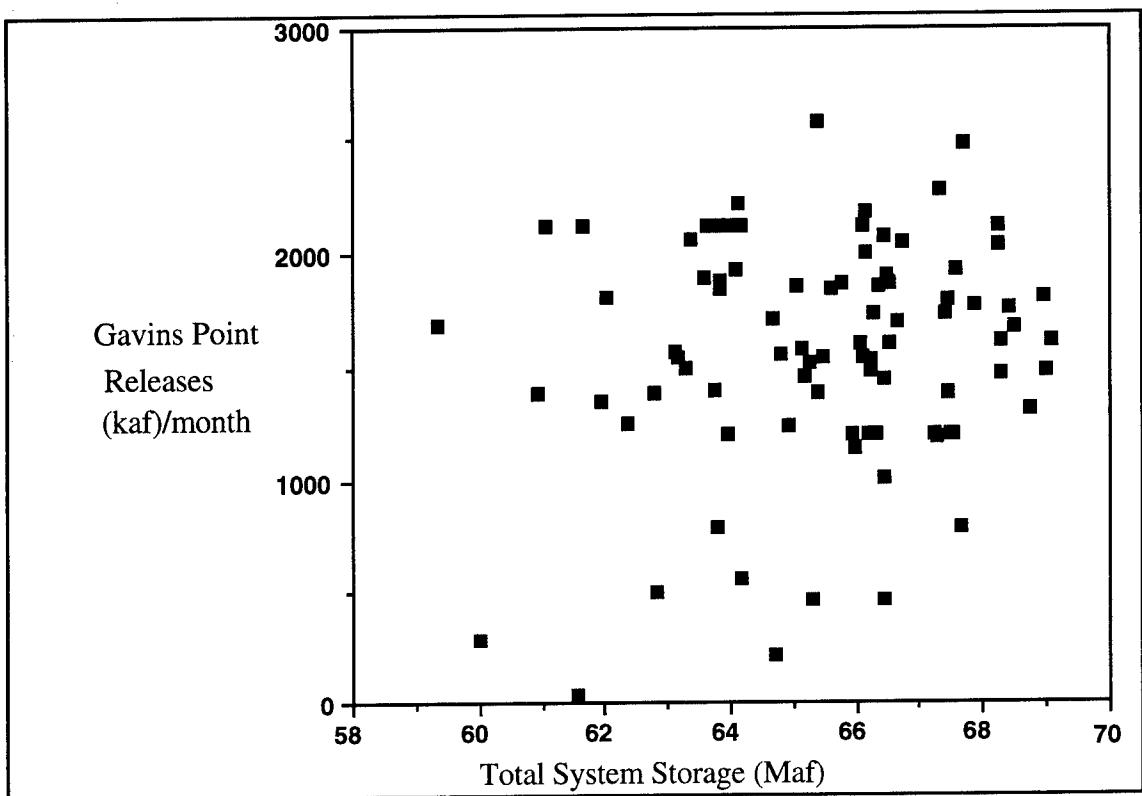


Figure 3.3: July Relationship between Systemwide Storage and System Release for the Missouri River System

Principal Components Analysis

For multi-reservoir systems, a more elaborate statistical technique, principal components analysis, has been used to aggregate and disaggregate reservoir operation results (Saad and Turgeon, 1988; Saad, et al., 1992). Here, optimal reservoir operations are found for a large number of deterministic optimizations for a large number of equally likely inflow sequences. Synthetic streamflows were used to generate these inflow hydrologies. The results of these deterministic optimizations are then analyzed by principal components analysis, finding the distribution of storage throughout the system for a given total system storage. The authors use this approach to reduce the number of reservoirs that need to be modeled for hydropower purposes, so that the problem is amenable to stochastic dynamic programming. As the authors note, the approach works only if there is considerable interdependence in the operations of the reservoirs.

Artificial Neural Networks (ANN)

A further statistical technique for inferring operating rules from deterministic optimization results is the application of artificial neural networks (ANN) to find and establish operating rules. Saad, et al. (1994) used ANNs to develop hydropower operating rules for a multi-reservoir system in Canada. The ANNs were trained using many sets of deterministic optimization results for the system, with each set being based on a different set of inflow hydrologies. For their application, the ANN method is used for disaggregating storage results from an aggregated stochastic dynamic programming representation of the system. The ANN approach would appear to be highly flexible and promising for general inference of rules from deterministic optimization results. However, the

calibration/training of an ANN can require a large amount of results data. For their example, the authors use 475 synthetic streamflow sequences, each 100 years long. This approach has been developed further incorporating fuzzy sets to speed convergence of the neural network calibration (Saad, et al., 1996).

Incremental Modification of Existing Rules

In many cases, deterministic optimization results are likely to agree rather well with those of existing operating policies. In such cases (USACE, 1995), the deterministic optimization results can be used to suggest areas where existing operating rules might be refined to make incremental improvements. Overall, in such cases, the results can be seen as providing rather rigorous support for the existing rule structure, since deterministic optimization results represent an upper bound on system performance. Agreement between optimized and current operating policies should not be surprising, given the often decades of experience, thought, and simulations which have typically gone into developing current operating policies.

Application of HEC-PRM to the Columbia River System provided results which were in relatively close agreement with existing operating policies, represented by the HYSSR simulation model (USACE, 1995). As summarized in Figure 3.4, HEC-PRM operations agreed closely with those of current operations (HYSSR) in terms of total system storage. HEC-PRM tended to distribute water in the system a little differently, however, keeping Grand Coulee reservoir full a greater proportion of the time.

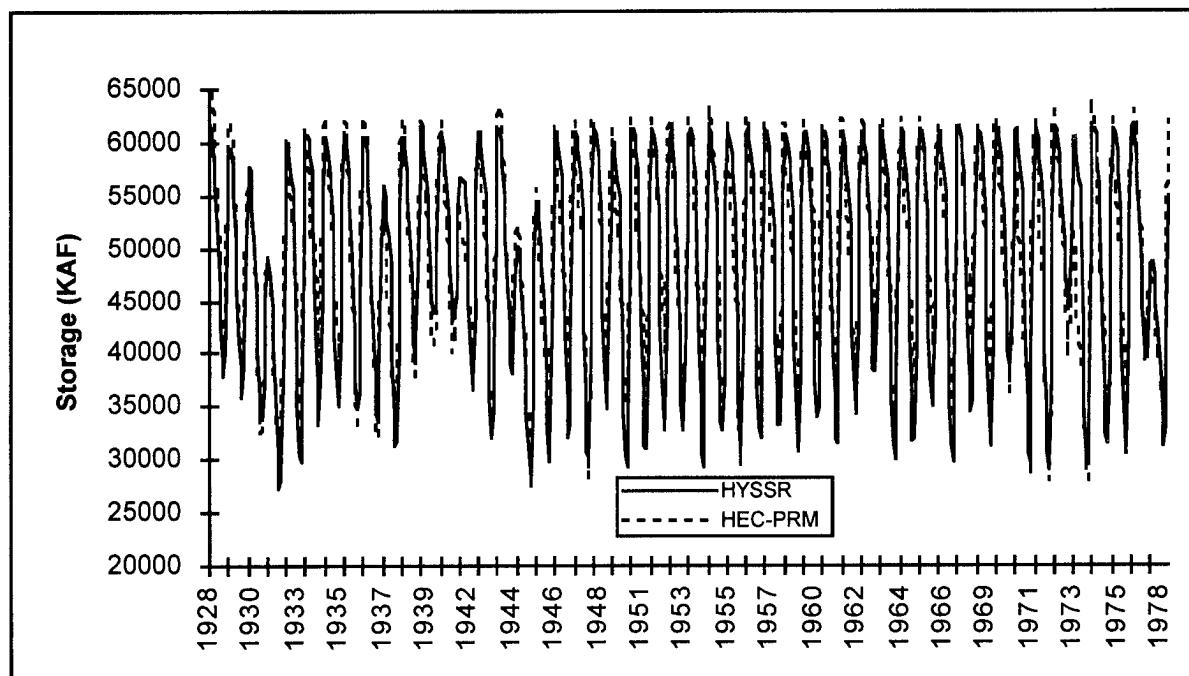


Figure 3.4: Comparison of Systemwide Storage Between HEC-PRM and HYSSR for the Columbia River System

3.4 Mixed Simulation-Optimization Approaches

Simulation-optimization approaches to developing operation rules for reservoirs employ optimization models to suggest initial operating rules and simulation models to test and refine these rules. This process may involve several cycles of optimization and simulation runs, often conducted in a fairly adaptable and flexible, but systematic way. Almost every practical rule-making exercise undertaken using optimization has conjunctively employed simulation modeling (for example: Jacoby and Loucks, 1972; Evenson and Moseley, 1970; King and Evenson, 1972; Toebes and Rukvichai, 1978; Bhaskar and Whitlatch, 1980; Karamouz, et al., 1992; USACE, 1994b; Lund and Ferreira, 1996). Some of the general rationale and uses for using simulation with optimization are presented in Table 3.1 and discussed below.

Rationale for Use of Simulation

There are several reasons to employ simulation in conjunction with optimization for reservoir rule-making. First, optimization models must typically be somewhat simpler than simulation models of a reservoir system. Optimization models typically require that definitions of the system and its objectives conform to specific mathematical conditions needed to implement a solution method. For HEC-PRM, an example is the requirement that all penalty functions be convex. Simulation models suffer much less from such limitations. This makes it possible to test rules developed from optimization results with more realistic simulation models. The greater realism of simulation models also provides opportunities to refine operation rules suggested by optimization results to make them more appropriate for the real reservoir system. The greater detail available to simulation models can allow the "optimized" operations to be improved in their details.

A second reason for employing simulation models in rule-making with optimization is the often greater ability of simulation modeling to perform "what if" studies. Specific flood control or drought scenarios can be studied easily using proposed operation rules in a simulation model. This would be awkward and often inappropriate for optimization models.

Third, simulation models typically run much faster than optimization models. A larger number of specific cases can be studied by simulation modeling than would be possible by optimization. However, optimization results might suggest some of the more fruitful scenarios to be tested.

Fourth, simulation models can avoid the false perfect foresight entailed in most deterministic optimization models. Rules can be tested and refined under more realistic conditions of future knowledge using simulation.

The final, and perhaps most important, reason to employ simulation models is the greater acceptance enjoyed by simulation modeling and the frequent relative ease of explaining simulation results. Even where operating rules are unchanged by simulation modeling, simulation modeling is probably necessary to render the rules understandable and acceptable to concerned technicians, decision-makers, and other interested parties.

Uses for Simulation

Rule Refinement

Since simulation models can both represent the reservoir system in greater detail and be executed more quickly than optimization models, simulation models are useful for refining the details of operation plans suggested by optimization results. As such, the optimization-based

suggested rules may serve mainly as a point of departure for more traditional simulation studies of operation plans (USACE, 1994b).

Simulation modeling also can be used to refine the optimization model (Karamouz, et al., 1992). In this case, a cycle of optimization, rule-making, and simulation modeling proceeds iteratively until a satisfactory set of rules is developed.

Rule Testing

Again, since simulation models can represent the system in more detail and have already gained some acceptance, in most cases, simulation modeling is a rather inexpensive and effective approach to testing operation plans developed from optimization results. Such simulation tests can answer several questions:

1. Do the suggested rules closely match the storage and release behavior from the optimization model? By implementing the suggested rules in a simulation model, rule-based storages and releases can be compared with those obtained directly from the optimization model (Bhaskar and Whitlatch, 1980; USACE, 1994b). This comparison can be used to see if the suggested operation rules well represent the optimization results.
2. Are the suggested rules feasible? Unless the suggested rules are thoroughly thought out, it can be possible for rules to suggest impossible behavior. For instance a release rule based solely on the month can suggest release volumes in excess of available storage and inflows.
3. Are the suggested rules really optimal? Since a simulation model can usually represent the reservoir system in greater detail than an optimization model, implementing the suggested operation rules in a simulation model and performing sensitivity analysis on the parameters in the suggested rules can conceivably improve the optimality of the suggested rules. A similar test is to compare the detailed performance measurements from a simulation model employing existing operation policies with those from a simulation employing the suggested operation plan. If the optimization model represents too great a simplification of the real system, existing operations might in fact be superior to those suggested by the optimization model.
4. Do the suggested rules perform well under extreme detailed scenarios? It is often desirable to test a proposed operation plan under detailed flood control, drought, or emergency operation circumstances. If the suggested operations are not suitable for such emergency operations, the suggested operations, the importance of the chosen scenarios, and other responses to the proposed scenarios might be further examined. Often, further optimization and simulation studies would be useful for such questions. For instance, the introduction of further constraints to the optimization to facilitate emergency operation can give cost estimates for preparedness for such emergencies, such as the cost to averaged annual operations of maintaining additional exclusive flood control storage. In some cases, there might be less expensive approaches for emergency preparedness.

Implementation Issues

The use of simulation in conjunction with optimization is greatly facilitated by the prevalence of existing simulation models for reservoir planning and operation studies. Almost all large reservoir systems have one or more existing simulation models. Still, most existing reservoir simulation models are likely to require considerable modifications to accept the diverse forms of operating rules that are likely to be developed from deterministic optimization (e.g., HEC-PRM) results.

In many cases, the most difficult aspect of simulation studies of this nature is the incorporation of more explicit economic or environmental performance indices in an existing simulation model. While adding explicit performance indices to simulation models may be burdensome and time-consuming, the presence of economic and environmental performance indices in a model can be of long-standing utility long after an operation plan study is completed (USACE, 1993b).

Table 3.1
Rationale and Uses for Simulation Modeling in Optimization Rule-Making

Rationale

Simulation models typically represent the system better than optimization models.

Simulation models perform some "what if" studies more easily than can optimization models.

Simulation models typically run faster than optimization models.

Simulation models avoid the perfect foresight of most deterministic optimization models.

Simulation modeling is typically better understood and accepted than optimization modeling.

Uses of Simulation

Refinement of suggested optimization-based rules to increase realism in system operation.

Testing of suggested optimization-based rules for:

- feasibility
- detailed operational implications
- comparison with existing operation plans
- evaluation of desirability using more detailed operational performance measures

3.5 Testing and Refining Operating Rules by Simulation

For long-term operating rule studies, simulation modeling is an essential complement to long-term deterministic optimization results. Simulation modeling studies provide the ultimate means of testing, refining, and justifying operating rule suggestions arising from optimization model results.

The process of testing and refining operating policies with simulation models begins with a formulation of the outline of an operating rule strategy, as inferred from optimization model results (USACE, 1992, 1994b; Lund and Ferreira, 1996). This outline can be partial or constitute a complete set of operating rules, depending on what could be inferred from the optimization model results.

A simulation model of the system must then be developed which allows this structure of long-term operating procedures to be examined. In some cases, such a model will already exist,

but where the operating rule strategy suggested by the optimization results does not conform with the structure of rules in the existing model, a new model often must be developed.

To begin testing, it can be useful to provide focused comparisons of particular components of the operating strategy suggested by the optimization model. These comparisons typically would be between the optimization results and simulation results for only those rule components being tested. For example, in testing fixed storage rules for several reservoirs on the Missouri River system, these fixed storage levels were compared to the slightly variable operations under the HEC-PRM model. Similarly, a storage allocation rule was tested by comparing the HEC-PRM storage results with those simulated using the storage allocation rule with releases determined by using total system releases from the HEC-PRM model (USACE, 1992, 1994b; Lund and Ferreira, 1996). Such partial testing of the operating strategy allows identification of weaker parts of the proposed strategy and provides justification for stronger parts of the strategy.

Refinement of the operating rules typically follows the testing of partial and complete preliminary operating strategy. These refinements naturally concentrate on the weaker parts of the strategy found in the earlier tests.

Final testing of the refined operating policies is through comparison with the original optimization results (they should agree mostly but rarely perfectly) and comparison with the current operating policy (for reasonableness and to ensure that current operations do not better match the optimization results). Ideally, the incorporation of detailed penalty functions into the simulation models of the system would allow more system performance based comparisons, such as the overall penalty of operations under current operation policies, suggested operation policies, and direct optimization model results. Where an overall penalty function scheme is unavailable, comparisons of narrower statistics of performance, such as hydropower production, flood damage, or water supply shortage, can be employed.

3.6 Some Pitfalls in Operating Rule Development

There are several potential pitfalls in the development of operation plans from deterministic optimization results. Most of these can be detected by the use of simulation studies to test and refine suggested operation plans. Some of these pitfalls are probably of mostly academic importance, but may have practical importance in specific cases.

Infeasible Operations

It is possible for the set of rules suggested by optimization results to result in infeasible operations. Infeasible operations are those that would not be allowed by the constraints in the original optimization model or not physically possible in the real reservoir system. The likelihood of infeasible operations increases when the reservoir system faces more severe drought or flood events than those present in the hydrology entered into the optimization model. Infeasible operations are also more likely to result from suggested rules which do not closely mimic the optimized operation of the reservoir system. An example of an infeasible operation is a rule which specifies releases greater than the sum of the available storage and inflow.

Technical Suboptimality from Failure to Represent Uncertainty

The results of the deterministic optimization model represent an ideal operating policy, with perfect forecasting of future inflows and perfect predictions of the value of different reservoir purposes. As such, it is unlikely that any set of rules triggered by current operator knowledge

(such as current month, storage, and inflows) will be able to perfectly mimic the optimized results. This implies that the suggested rules will not produce as good an operation as that given directly by the optimization results. The deterministically optimal operation is an upper bound, in this sense, on overall possible system performance.

The divergence between the rules suggested by the optimization results and the optimization results represents, in some sense, the cost of uncertainty in streamflow forecasts. It may be possible for a more rigorous stochastic optimization to provide rules for which this divergence would be less. However, such stochastic optimization is difficult or impossible for many real reservoir operation problems (Miquel and Roche, 1983).

Technical Suboptimality from Optimization Model Simplification

As mentioned before, most optimization models require some simplification of the real reservoir operation problem. For HEC-PRM the need for the objective function to be convex is such a simplification. This implies that the optimal operations suggested by the optimization model may not be the real optimal operation. While some of this phenomena may be tested by simulation modeling, the exact optimal operation for the real system is in practice usually unknowable.

Oversimplification of Rule Forms

There is a great temptation to seek a few simple rule forms when developing operation plans from optimization results. This principle of parsimony is generally very useful and well accepted in professional and scientific fields. However, it may be possible for more complex rule forms to more closely mimic the optimization results and improve reservoir operations.

Overly Complex Rule Forms

Rule forms that are overly complex might more closely mimic the results of the optimization model. However, too complex a set of operating rules can result in a degree of spurious correlation between rule-based operation and optimization results. Complex operating rules also make simulation studies more difficult.

Replication of Existing Operation Through Rule Form Selection

If current operation plans are used as a guide for developing new operation plans from optimization results, it is likely that the "new" operation plans will be very similar to the existing operation plan. The use of the same form for new operating rules as existing rules can result in a close replication of existing policies. Some attempt should always be made to see if rule forms different than existing forms can closely mimic optimization results. Despite such efforts, in many cases it is likely that "optimal" operating plans will be rather close to existing operating plans (USACE, 1995).

Absence of Stable Operating Patterns

It is possible that the optimization results might show no consistent pattern of operations. Such a case might result where operating conditions in each year are driven by penalties from a different operating purpose, resulting in qualitatively different operations for different years. Results also may show little pattern if annual changes in hydrology drive the system to very different operations with nearly the same penalty value. Thus, a consistent pattern of near-optimal operations can be obscured by purely optimal operating decisions with perfect foresight. Fortunately, while optimization results always lack perfect clarity, there is usually a fairly clear pattern to much of the optimized operation decisions.

Several Operating Strategies Yield Similarly Good Results

There is a possibility that more than one set of reasonable operating rules could produce very similar reservoir operations, all similar to the optimization results. This is a common trait of "inverse problems", such as this, where the results are known, and an underlying process must be inferred. Several processes, or sets of rules, might yield the same, or nearly the same, results. The major controls on this problem are the use of long series for hydrologic input and, when given a choice, the preference for simpler operating rules. The long series for hydrologic input produces large amounts of data which should reduce the likelihood of several good rule sets being found. Where several sets of operating rules are found which produce equally good results, the "simpler" forms, or those most simply employed or explained, would generally be chosen ("Occam's razor"). Again, simulation modeling should help in addressing this problem, were it to arise.

3.7 Conclusions

The development of strategic operation plans from deterministic optimization results using long hydrologic records has advantages (and disadvantages) over traditional approaches employing simulation and engineering judgment or stochastic optimization. This approach to operation plan development has a long history in the engineering literature with a large number of plan development approaches being suggested. These rule development approaches are summarized in Table 3.2.

Table 3.2: Summary of Approaches for Developing Long-Term Operating Policies from Optimization Model Results

Rule Development Approach	Summary Comments
Intuitive Approaches	Probably of some use in all cases.
Operation rule theory	Useful for developing operation strategy, but probably insufficient alone.
Regression	Often difficult for real applications, but often can identify important variables.
Principal components analysis	May work well in some cases.
Artificial neural networks (ANN)	Requires unusually large amounts of streamflow data.
Incremental modification of existing rules	This should be the most promising approach, given good existing rules, small changes in operating purposes, and an existing simulation model.
Mixed simulation-optimization	Probably essential in most cases.

The selection of a representative hydrologic input for the optimization model is important and can vary with the narrow purpose of the optimization study. If the application of optimization is to narrowly focus on operations during wet, dry, or normal years, then different sets of design hydrologies would be selected. If the application is to assess desirable operating rules over a wide range of conditions, then long-term optimization based on long historical, extended historical, or synthetic streamflows would be desired.

In general, a combination of a variety of plan development approaches is likely to be preferred. In particular, the use of simulation modeling in conjunction with optimization results is almost essential to the technical and practical success of any rule-making exercise based on optimization results.

Chapter 4

Seasonal Rules from Optimization Studies

4.1 Introduction

This chapter reviews approaches for developing seasonal operating rules from the results of HEC-PRM or other deterministic optimization models. Issues discussed include those of setting up seasonal operating problems for a HEC-PRM application, different approaches to making sets of seasonal runs, and the interpretation of seasonal results to offer advice for seasonal operations. The approach and many of the methods discussed here have been applied preliminary to the Columbia River System (USACE, 1996). A summary of this work appears in Appendix B.

The development of seasonal operating rules using deterministic optimization modeling begins with several decisions regarding the optimization model. These decisions include:

1. establishing penalty functions,
2. setting initial conditions,
3. how to handle end-of-period storage targets and penalties, and
4. the selection of relevant hydrologic inputs.

The various objectives that can be used in the optimization model are presented in Chapter 1, and are substantially similar for both seasonal and long-term operations studies. However, as discussed below, these objectives can be somewhat different in a seasonal operations context.

Seasonal operations studies offer an opportunity to provide more realistic and specific advice than the long-term studies discussed in Chapter 3. For seasonal studies, the initial conditions are well specified by current storage levels. Penalty functions for operations are usually better known for the short term than for the long term. And there is often inflow forecast information which can be used to improve our estimates of future inflows, beyond the hydrologies available for long-term studies. However, seasonal studies, being of short duration, tend to require greater attention to end-of-period storage targets and penalties.

Beyond these details of setting up seasonal operations optimization studies, this chapter also discusses some aspects of how these modeling results could be interpreted and incorporated into seasonal operations decision-making. These seasonal decision-making uses of optimization model results include:

1. annual and seasonal operation plan development,
2. refill reliability assessment,
3. yield reliability analysis,
4. drawdown and refill ordering advice,
5. establishment of short-term operating targets, and
6. external justification (or a second opinion) of current operating policies.

4.2 Penalty Functions and Initial and End-of-period Conditions

Before a seasonal operation run of HEC-PRM or any other deterministic optimization model can be completed, some specific formulation problems must be addressed. These include the penalty functions for general operations, setting the initial storage conditions, and specifying any end-of-period conditions, targets, or penalties that should apply.

Penalty Functions for Seasonal Operations Studies

The penalty functions for seasonal operation planning studies can be a little different from those appropriate for strategic operating rule studies. Seasonal studies are made typically for operations during the upcoming operating season or year. In this specific time-frame, the appropriate penalty functions for system operation are not necessarily those of some future expected level of development. For the coming year, or season, the appropriate set of penalty functions are those of current development levels, perhaps modified for particular maintenance or other activities planned for the coming year.

Initial Conditions

The initial conditions for seasonal operations studies should represent current storage. It is important that actual current storages be used for the initial conditions, since seasonal operation studies are typically for short periods, with the first few periods being of the most immediate interest. Thus, there is little or no ability to allow the system to "warm-up," as discussed in Chapter 3, without losing important operational results.

End-of-Period Conditions

There are several approaches to handling end-of-period storage targets for seasonal operation studies. These approaches each have different advantages and disadvantages.

Fixed Storage Constraints

The simplest approach to handling end-of-period storages is to firmly establish them with constraints. In HEC-PRM, end-of-period storages can be constrained to be above some minimum storage level and/or constrained to be below some given storage level. The end conditions can also be set exactly to specific storage levels. While this constraint approach is simple to implement, in cases where final storage targets are actually "soft", the constraint approach can lead to sub-optimal operation results from the model in attempting to achieve inflexible constraints.

In a modification of this approach, the fixed storage targets for the last period might be varied, depending on whether the year is wet, normal, dry, or based on some other condition known to the system operators. Varying carryover storage based on year type is not uncommon for some reservoir systems and represents a form of hedging, so this approach can be both reasonable and a reasonable representation of current practice.

Storage Targets with Penalties

Special end-of-period penalty functions can often better represent the desirability of achieving particular end-of-period storage targets. Such penalties apply to storage only in the last time-period of the analysis. Figure 4.1 below illustrates how such penalty functions can be constructed.

The end-of-period storage penalty function is anchored with a target storage level, the ideal level of storage for the system at the end of the season or year. However, missing the target might be desirable if this sufficiently improves operation of the system during the analysis period. This flexibility is allowed by providing penalties for missing the target, allowing the optimization model to calculate the trade-offs of satisfying the target versus improving operations during the analysis period.

For systems where the target storage is established to provide carryover storage for the coming year or season, then the penalty for storing water less than the target should be the value of

carryover storage, as shown in Figure 4.1. Where the target storage is set to provide flood control space for the coming season, then the penalty for exceeding the target storage should represent the expected value of flood damages avoided by retaining that amount of flood control space in the reservoir.

The storage targets with penalties approach was used in preliminary seasonal studies of the Columbia River system (USACE, 1995). In this study of the system's refill season, the value of energy storage was used to estimate the slope of the end-of-period penalty function. Thus reservoirs higher in the watershed had higher penalties for failing to achieve the target. This penalty-based approach to the last period thus allows the model to decide where storage should be kept during excessively dry or wet years, and how to trade off operational benefits within the analysis period against benefits in later periods.

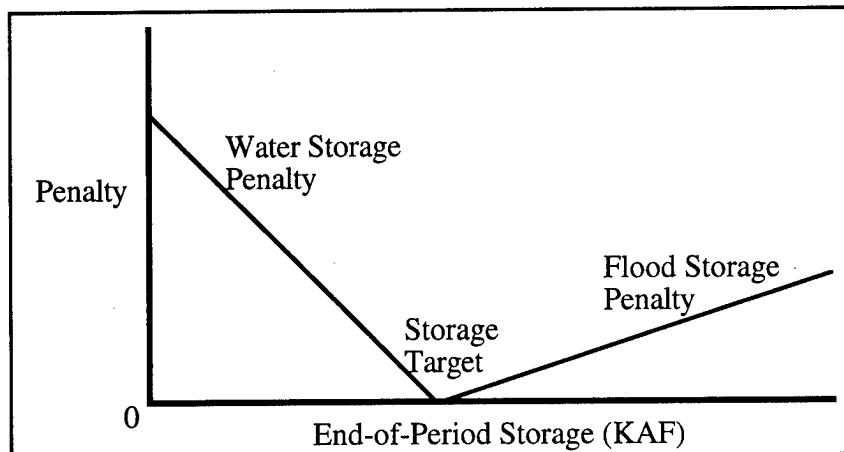


Figure 4.1: End-of-period Storage Penalties

Additional Cool-off Period

Another approach to the end-of-period problem is to merely extend the analysis period into the future long enough to dampen end-of-period effects, and then discard the results for this extension. This approach essentially provides a "cool-off" period for the analysis. This approach is simple to implement, but has the following problems:

1. Computation time is greater for the longer model runs.
2. Peculiarities of the particular inflows during the cooling-off period can substantially affect model results during the period of concern.

The use of a long cool-off period can be desirable if all of the following conditions are met: 1) the system has very long streamflow records or an ability to generate reliable statistical synthetic streamflows, 2) system storage capacity is small relative to average inflows, and 3) there is significant cross-seasonal or inter-annual correlation in streamflows. Under these conditions, an extended cool-off period might be a practical method for improving representation of likely future inflows to the system. This is an instance where decisions on handling end-of-period conditions and hydrology selection can interact.

4.3 Hydrologic Input

There are two basic approaches to establishing the hydrologic inflow records for seasonal deterministic optimization model runs. Commonly, seasonal operations are tailored to respond to one or a few potential hydrologic or streamflow events, sometimes referred to as design events. The selection of these design hydrologic events implied some particular concerns for system operations. A second approach to seasonal operations studies employs many potential hydrologic events to provide a fuller and more representative picture of how the system might or should operate in the coming season(s).

"Design Event" Hydrologies

The first approach is to select one or a few "design" hydrologies to be used for the model. As in the case of long-term studies discussed in the previous chapter, these "design" hydrologies can be chosen to illustrate particular inflow conditions of concern, such as floods, droughts, or typical years. For seasonal operation studies intended for use near the time of actual operations, forecast flows also might be available for use.

Single Forecast Event

Forecast flows are often developed for many large watersheds. These forecast flows are commonly used in simulation modeling to evaluate alternative seasonal system operations. For seasonal studies, these forecasts are particularly useful for systems with a large proportion of snowmelt runoff, where the amount and distribution of snowpack can influence runoff for several months. Even in this case, the snowpack melting rate and the possibility of later precipitation can lead to large errors in single forecasts. (In many cases, several forecasts are made, assuming different melting rates and future precipitation conditions.) Nevertheless, such forecasts are often readily available and offer a seasonally-varying basis for estimating streamflows for modeling studies of seasonal operation.

"Typical" Year

The use of a "typical" year of hydrologic record for seasonal operating studies is discouraged. A "typical" year's hydrology might give some qualitative idea of how a system might desirably operate during such a specific year. However, such a model run provides little idea as to how sensitive such operations are to hydrologic variability or any important range of hydrologic conditions.

"Critical Period"

For seasonal operations concerned with drought, it is sometimes useful to examine how the system should operate if faced with a repeat of the drought of record, or severe droughts from the historical record. Thus, the "critical period" is often used as the hydrology for seasonal operation modeling studies, particularly those used to estimate firm water or power yield for contracting purposes for the coming year. An optimization model's advice for operation for the coming season assuming a "critical period" hydrology is likely to be rather conservative.

"Design Flood"

Similarly, it can be useful to see how well the system would operate during the coming year of season if very wet, or flood, hydrologies are assumed. It is often desired, particularly early in a season, to maintain flexibility to manage the great floods of record or other design floods. The results of an optimization model, using design floods for hydrology, provide some

understanding of the current system's ability to manage such floods and produce other benefits under such wet scenarios.

Multiple "Design Events"

In many cases it is likely to be desirable to solve the optimization model several times, under different hydrologic "design events." Such a strategy might solve the optimization model for several inflow forecast scenarios, and perhaps the "critical period" and one or more "design floods". Where optimized operations agree closely under a wide range hydrologic scenarios, it is likely that such operations are desirable. Where results diverge, they may offer some insight into how desirable operating decisions should vary as hydrologic conditions unfold during the coming season. More elaborate versions of this multiple event approach are discussed below.

Many-Event Hydrologies

System inflows for the coming season or year are typically uncertain. The use of many inflow hydrologies is appropriate if it is desirable to examine the ideal operation of the system over a wide range of representative events. There are several approaches to selecting such inflow hydrologies (Hirsch, 1981b).

Repeats of Historical Seasons

For seasonal operation studies, a historical record of n years (corrected to unimpaired flows) can be used directly to provide n representative scenarios of future streamflows for the coming year. This approach was taken by Hirsch (1978) using simulation modeling for water supply risk studies. This approach is advantageous in climates where intermediate-term forecasts have little value (and thus the historic realizations are just as good), or where there is insufficient hydrologic expertise or budget to create forecast flows.

Multiple Forecasts

For some large systems, multiple seasonal inflow forecasts are made. Perhaps the largest-scale creation of inflow forecasts is done by for the Columbia River System. During the system's variable drawdown and refill seasons (January-July), unimpaired flows from the historical record (65 years) are modified to reflect current snowpack and watershed conditions. This updating is repeated monthly throughout the refill season. The existence of these many forecasts, updated monthly, is almost ideal for seasonal operations studies. These inflow forecasts have been used as inputs for HEC-PRM application to the Columbia River system for seasonal operation (USACE, 1995).

Synthetic Streamflow Forecasts

Various forms of synthetic streamflow forecasts also can be used to provide input hydrology events for seasonal operations studies (Hirsch, 1981b). Hirsch (1981a) combined direct use of historical data and synthetic hydrology methods for risk analysis of the New York City water supply reservoir system. Each historical seasonal streamflow trace was modified for the statistical correlation of present flows with future inflows. This should improve the representativeness of the inflow record by adjusting it for the information contained in current flows.

4.4 Operational Uses of Many-Event Runs

The results of many-event seasonal operation runs from a deterministic optimization model can be of more direct and immediate use than the results of long-term optimization runs discussed in Chapter 3. Seasonal operation runs all begin from the same current storage levels, with hydrologies which should be representative of near-future inflows. The results of such seasonal optimization runs can be displayed and employed in several ways.

There are several ways that the results of seasonal deterministic optimization models can be useful for seasonal operations. These include relatively direct uses of results for near-term operating decisions, estimating refill probabilities, and the ordering of refill and drawdown for multi-reservoir systems. Examples of these techniques can be found in USACE (1996), summarized in Appendix A. The development or adjustment of operating rule curves or other operating policies for the coming year or season is an additional use, discussed in the section following this one.

Traces of Storages and Flows

The presentation of many-event runs can often be quickly and simply displayed through "position analysis" plots (Hirsch, 1978, 1981b). Such plots show a trajectory for each optimization (or simulation) run, each run representing a particular assumed future hydrology. The pattern and relative frequency of storage or flow for each time-step can be useful for seasonal operation planning. In the plot below, Figure 4.2, for ten representative hydrologies, all storages begin in the first month at the initial storage (5,000 KAF). Except for the driest scenarios, there is some refill, followed by drawdown until the 5th month to the 7th month, followed by often rapid refill to 10,000 KAF in all cases.

Where many hydrologic series have been used, position analysis plots can become cluttered. In such cases, quartile or percentile plots, like those described in Chapter 3, can be used to clarify the behavior of the system. Figure 4.3 is a quartile plot of the results appearing in Figure 4.2, showing maximum, 75th-percentile, median, 25th-percentile, and minimum storage values for each month of the coming year. These results can be used for both reservoir management and reservoir operation purposes on a seasonal basis.

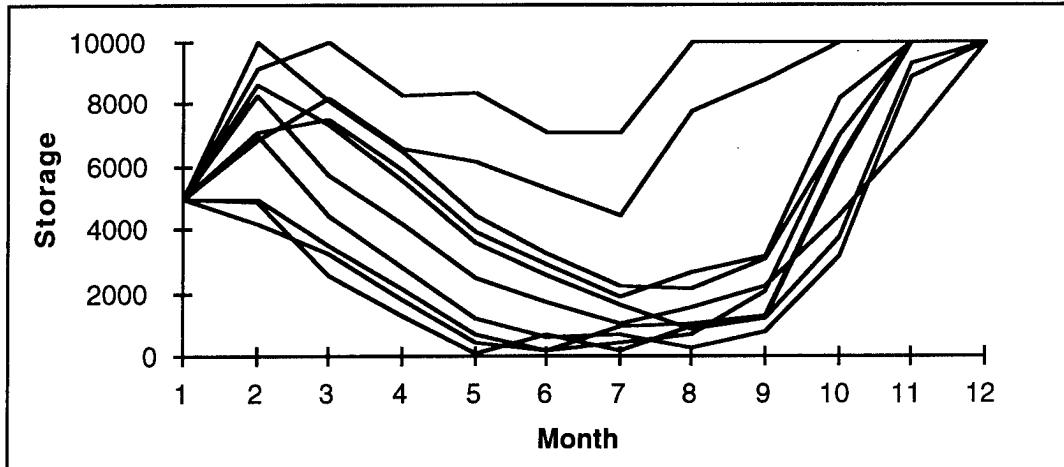


Figure 4.2: Example Storage Position Analysis Plot

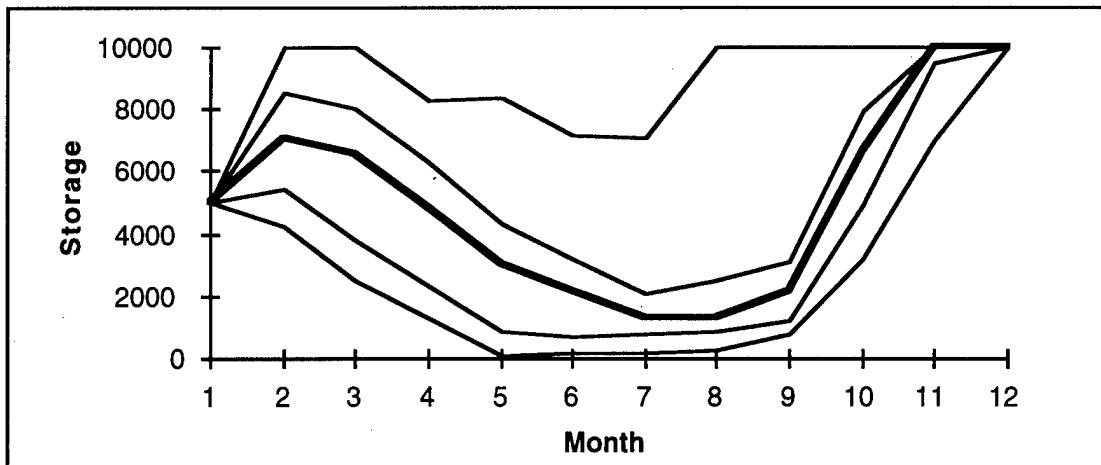


Figure 4.3: Quartile Plot of Storages from Position Analysis Plot in Figure 4.2

Suggestions for Near-Term Decisions

The results of deterministic optimization models for seasonal operation problems sometimes suggest direct operating policies, in the form of release or storage targets for near-term operation decisions. The example in Figure 4.4 illustrates this point with the position analysis plot of Mica Reservoir storage (on the Columbia River system) for the January-July, 1994 period (USACE, 1995). For 48 forecast hydrologies (representing historical inflow hydrologies updated for January 1 snowpack and forecast information), the storage target advice is rather consistent for January - April, 1994.

The results for Arrow Reservoir, from the same Columbia River study, appear in the position analysis plot in Figure 4.5 and its quartile plot cousin in Figure 4.6. Here, the advice is not nearly so precise, but is still fairly consistent, at least in qualitative terms. In the first few months of the operating season (January-March), the reservoir should be drawn down, and there should be a strong basis for the reservoir refilling, or coming close to refilling, by the end of July. The percentiles of storages for various near-future months should provide a reasonable basis for assessing a reasonable rate of draw-down.

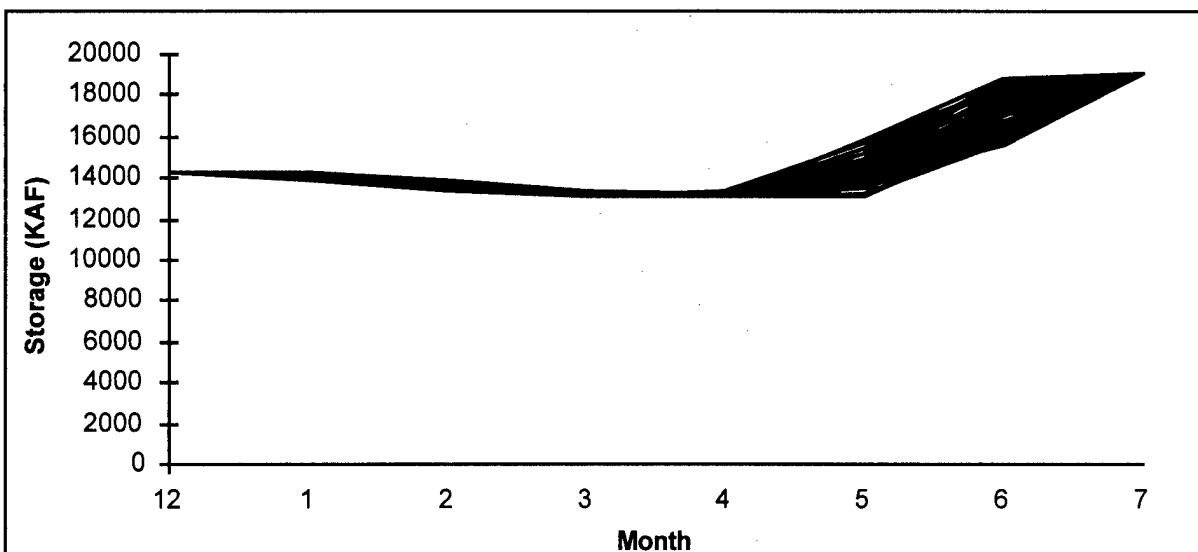


Figure 4.4: Example Position Analysis Results for Mica Reservoir, Columbia R. System, for January-July, 1994

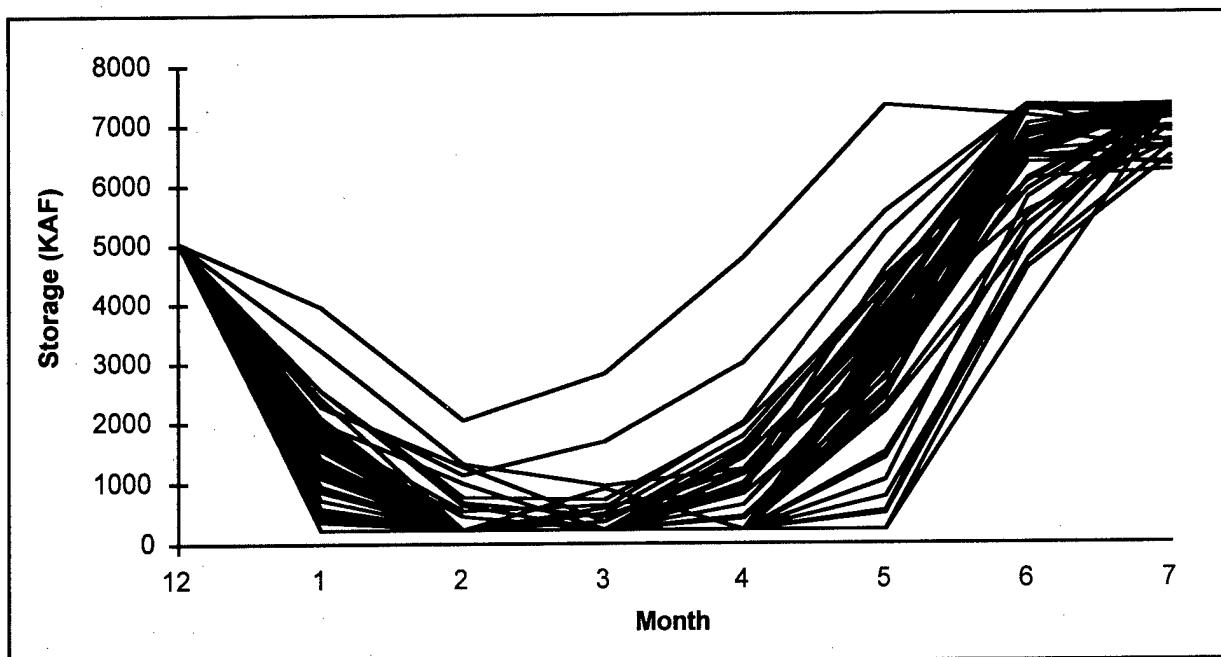


Figure 4.5: Example Position Analysis Results for Arrow Reservoir, Columbia R. System, for January-July, 1994

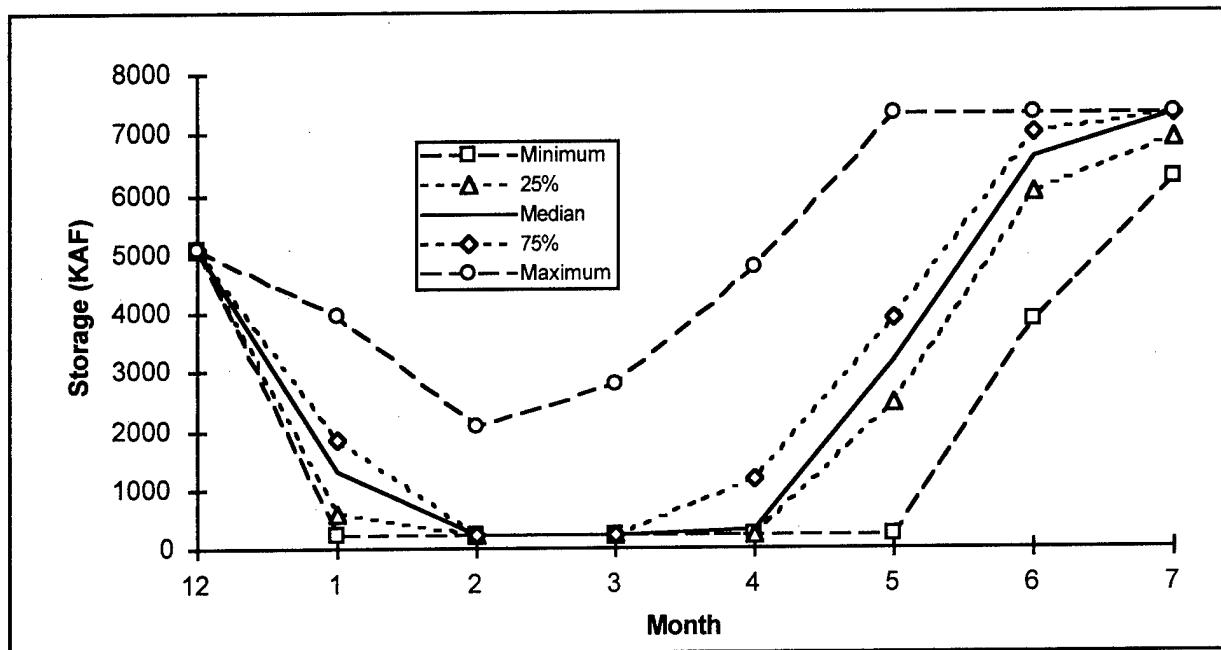


Figure 4.6: Example Storage Quartile Results for Arrow Reservoir, Columbia R. System, for January-July, 1994

Refill Probabilities

Another result of using many-event hydrologies is an explicit estimate of re-fill or drawdown probabilities for the end of each season. This is obtained from the many estimates of maximum drawdown or refill obtained from the different hydrologic events. Similarly, these results can be used to estimate the month of maximum drawdown or refill. In the case depicted in Figure 4.3 above, refill probability for the ten inflow scenarios is 100%. For Figure 4.5 and 4.6, the refill probability is about 80%. Such refill probabilities can be estimated fairly closely using tabular results from the many-event runs, dividing the number of runs refilling by the total number of runs.

Order of Drawdown and Refill

For multi-reservoir systems, deterministic optimization model results can often be used to suggest a desirable ordering of reservoir drawdown or refill. Figure 4.7 shows a storage allocation plot for three reservoirs in the Columbia River System, using the January-March results from a HEC-PRM model for January-July, 1995 (USACE, 1996). For this application, the model was run using 47 historical flow sets, each updated to represent January 1 snowpack and watershed conditions. January-March is this system's "variable drawdown" season. The storage allocation plot in Figure 4.7 for a drawdown season, such as this, would be read from the full (or right-hand) side, drawing down each reservoir as shown as total storage decreases. The results in Figure 4.7 show a clear preference for drawing Arrow down almost entirely before drawing down Libby, with Duncan (having a capacity over 1 MAF) being drawn down largely before January, but emptied completely rather soon.

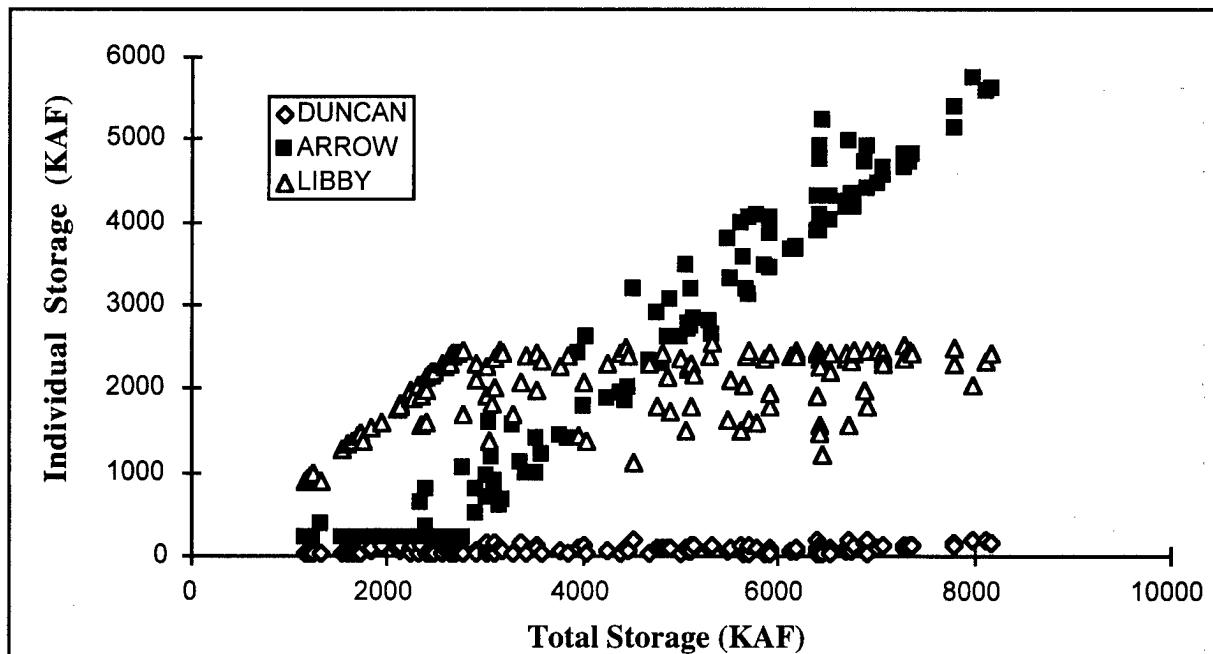


Figure 4.7: Storage Allocation Plot for Variable Drawdown Season, January-March, 1995 for Arrow, Duncan, and Libby Reservoirs in the Columbia River System

4.5 Development of Annual Operating Rule Curves

This section introduces the development of specific seasonal operating policies from many-event runs of deterministic optimization models for annual operation studies. This is similar to the problem addressed in Chapter 3, developing operating rules from long-term optimization results. Table 4.1 lists approaches available for modifying annual operating rules using results from seasonal optimization studies and provides some summary comments on their application to seasonal operation problems. Most of these approaches were discussed in detail in Chapter 3.

Table 4.1: Summary of Approaches for Developing Annual Operating Rules from Optimization Model Results

Rule Development Approach	Summary Comments
Direct use of results	Works well where optimization results agree for a wide range of hydrologic conditions.
Intuitive Approaches	Probably of some use in all cases.
Operation rule theory	Probably insufficient alone for detailed seasonal operations of real systems.
Regression	Often difficult for real applications; may have potential when existing or long-term rule forms provide the functional forms for regression.
Principal components analysis	May work well in some cases.
Artificial neural networks (ANN)	Requires unusually large amounts of streamflow data.
Incremental modification of existing or long-term rules	This should be the most promising approach, given good existing rules, considerable operating experience, and an existing simulation model.
Mixed simulation-optimization	Probably essential in most cases.

Developing operating rules in an annual or seasonal operation time frame, typically requires that the entire study must be completed within a few weeks or less. This is not nearly enough time for wholesale development of new seasonal operating rules each year or operating season. Thus, it usually will be best to develop seasonal operation studies which merely modify strategic operating policies established through experience and long-term reservoir operations studies. Of the operating rule development approaches listed in Table 4.2, intuitive approaches, reservoir operation theory, regression, principal components analysis, and artificial neural networks are probably of limited use for seasonal operation studies, especially if used alone.

Intuitive and reservoir operation theory approaches typically provide general advice on the desirable structure for system operating rules, but require a great deal of time to further develop, refine, and test. This makes them unsuitable, particularly when used alone, for seasonal operation studies, where such time is unavailable. Regression approaches have not been generally successful for developing operating rules for real multi-reservoir multi-purpose reservoir systems. However, where a generally successful rule pattern has been found, regression might be valuable for "re-calibrating" such established rules for more specific seasonal conditions. Principal components analysis might be similarly useful for particular systems where this approach appears to work well. Artificial neural networks also may be promising for some systems, but will typically require more hydrologic data than is usually available from historical streamflow records,

and therefore is likely to require synthetic streamflow generation. While these approaches are likely to have limited application for seasonal operating problems, they might, as discussed, be useful as part of hybrid approaches focusing on incremental modifications of existing operating policies. Where this incremental approach appears to work (i.e., the optimization model results tend to agree with general operating policies), the use of an optimization model also provides fairly explicit justification for the basic operating policies.

The most promising approaches would probably be incremental modifications of existing or long-term rules, mixed simulation-optimization, and occasionally, the direct use of optimization results. Incremental modifications to existing long-term rules is attractive because such studies should be able to be completed within a short seasonal operations time frame. Such studies would almost inevitably employ simulation modeling, making them also part of a mixed simulation-optimization approach. As described above, the direct use of optimization results would be attractive where the release decisions for a wide range of potential inflow hydrologies are in close agreement, as in Figure 4.4. However, even in this case, it is likely that this advice should be tested using more detailed simulation models.

The development of seasonal operating rules from deterministic optimization results is subject to many of the same pitfalls discussed in Chapter 3. However, again these problems typically should be controllable through use of simulation modeling.

4.6 Testing and Refining Operating Rules by Simulation

Because optimization models typically must represent reservoir operation problems in a somewhat simplified way, the operating advice suggested by optimization model results should almost always be tested, and probably refined, using more detailed simulation models. Just as with application of deterministic optimization to long-term operation problems, the results of deterministic optimization models to seasonal operations problems should also be viewed as a promising point of departure for more detailed optimization studies.

This use of optimization model results as a starting point for simulation studies can potentially reduce the number of simulation model runs required. If optimization model results can be used to focus simulation model studies on promising operations, it is likely that this combined use of optimization and simulation might require less overall modeling effort and time to complete, or allow for more detailed modeling studies than would be possible with simulation modeling alone.

Simulation model testing and refinement of rules for seasonal studies is a little different than for long-term operation studies. First, both the optimization and simulation studies must be completed in relatively short time-frame, often as short as several weeks, and perhaps several times during the year. Second, the objectives of operation should be somewhat clearer for near-term operations, reducing the need for sensitivity studies on the model penalty functions. Third, in some cases, such as where there is a significant snow-pack or soil moisture component to the annual hydrology, the appropriate representative hydrologies often can be of a narrower range than for long-term operation studies.

It is unlikely that optimization modeling will be anything more than a companion to existing simulation models for seasonal operation studies. However, the addition of optimization modeling to seasonal operation studies has some potential for improving system operations and for providing a more economically rigorous justification for operating policies.

4.7 Seasonal Operation Update Studies

The ability to complete seasonal optimization runs is quick order, with accompanying simulation refinement and testing, allows for the incremental updating of seasonal or annual operation plans throughout the operating seasons. This approach was illustrated for the Columbia River System (USACE, 1996). Here, it was found that the re-running of the HEC-PRM model, with updated current reservoir storage conditions and updated inflow forecasts, could improve the operational guidance offered for the system, compared to an initial run at the beginning of the season.

Where optimization models, such as HEC-PRM, can be used to speed seasonal system operation studies, they increase the ability to perform operation policy update studies throughout the operating season, and perhaps increase the usefulness of such update studies.

4.8 Management Uses of Many-Event Results

Water resources management includes a broad range of activities beyond reservoir operations. Management activities include contracting for power and water (or storage) supply, managing recreation facilities at reservoirs and downstream locations, and preparation for floods and water shortages for various purposes (navigation, water supply, etc.). This section presents a few examples of how the results of seasonal deterministic optimization modeling can be useful for water resource management activities that extend beyond reservoir operation. In a seasonal or annual time frame, these results should be most useful for contracting purposes and preparation for potential conflicts among reservoir purposes.

Probabilities of Storage Levels and Flows

Figure 4.8 shows an exceedence probability plot of storage for each month of results. Such exceedence probability plots can be generated for storages, flows, and diversions. For seasonal planning purposes such plots can indicate the probability of attaining particular levels of performance for reservoir system uses if the system were operated "optimally". Since deterministic optimization results assume that operations are optimized given perfect flow forecasts for the season, these cumulative distributions are likely to be a little optimistic, relative to actual operations, for many purposes. Table 4.2 identifies some indicators of reservoir system performance for several reservoir system purposes.

For example, the storage exceedence probability plot in Figure 4.8 could have several uses. Recreational planners can use the plot to estimate the likelihood of needing to extend or close a boat ramp which reaches only to where reservoir storage is 2,000 KAF, in which months closure would likely be necessary, and how likely it would be that an extension of a given length would be needed. Navigators using the reservoir would be able to estimate the likelihood of having various control depths into ports located on the reservoir, and in which months use of various ports would be most risky. In the case of Figure 4.4, if boat ramps extended to an elevation corresponding to 500 KAF of storage, there would be a 20% chance that the boat ramps would be unusable in month 6.

Table 4.2: Indicators of Performance For Various Reservoir System Purposes

Indicator	Purposes
Storage	Reservoir recreation Reservoir navigation
Flow	Downstream recreation Downstream navigation Flood Control
Diversion	Water Supply

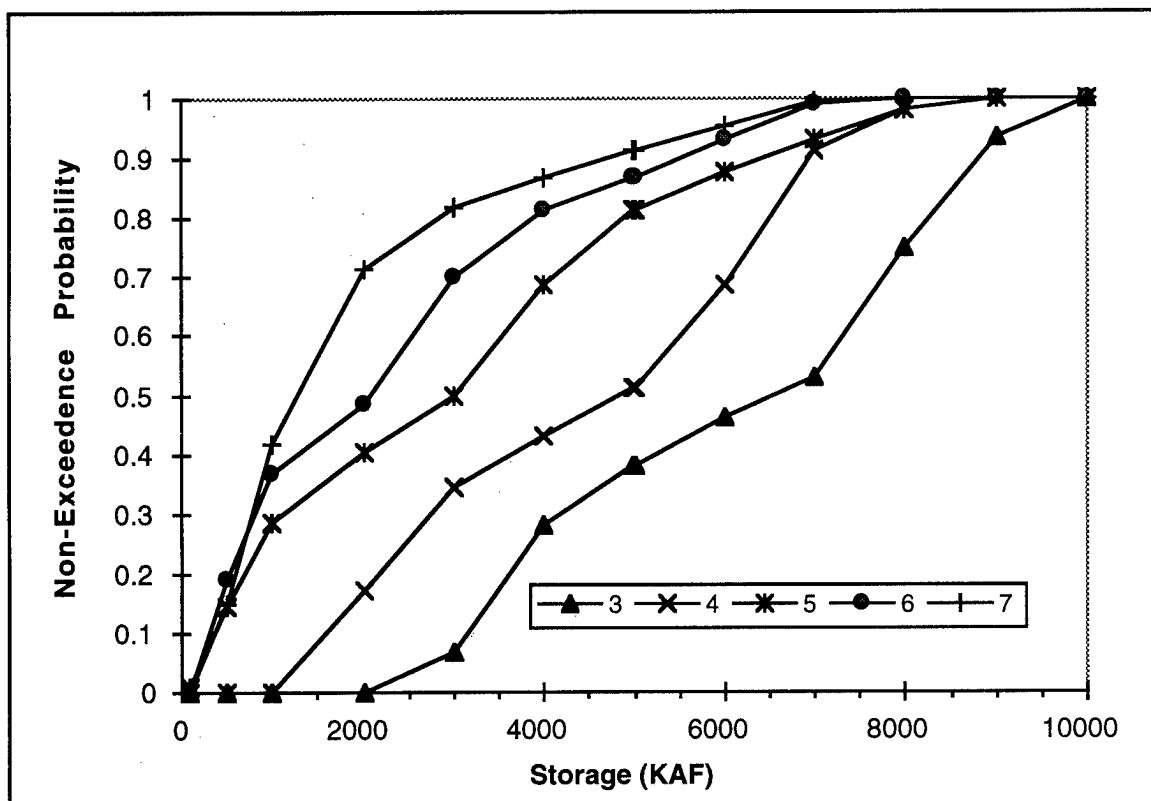


Figure 4.8: Monthly Exceedence Probability Plots for Storage for March through July

Water and Power Yields

The flow and storage results from an optimization model can be used to produce estimates of system water and power yields for each time-step during the period of analysis. Where multiple hydrologies are used, the results from these runs can be used to provide a distribution of water and power yields. Such results should be of some use in contracting for water or power provision during the coming seasons. Deriving such yield results using an optimization model is likely to provide somewhat optimistic estimates of yield, since it assumes perfect foresight and operations within each inflow scenario.

Figure 4.9 shows an example of a non-exceedence probability distribution for water yield for three different months. In this case, the plot, like many water yield plots, is rather jagged, reflecting the relative infrequency of shortage events where yields are less than target yields used for reservoir operations. This example is based on ten years of data.

Figure 4.10 is an example non-exceedence probability plot for energy yield for several months. This result required post-processing of the model results, using model storage and flow results as inputs to the hydropower equation. Such plots might be useful for the contracting of various forms of firm and non-firm power on a monthly basis for an upcoming year or season.

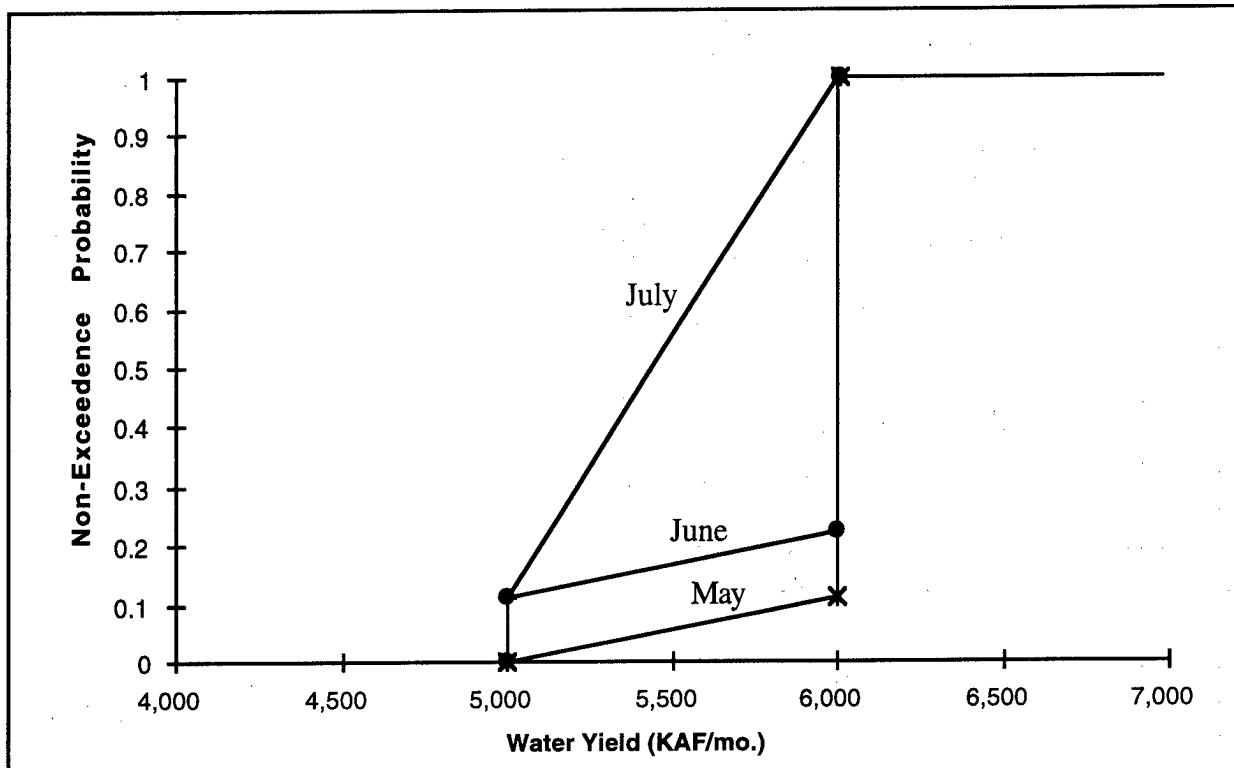


Figure 4.9: Example Non-Exceedence Probability Plots for Water Yield, May-July

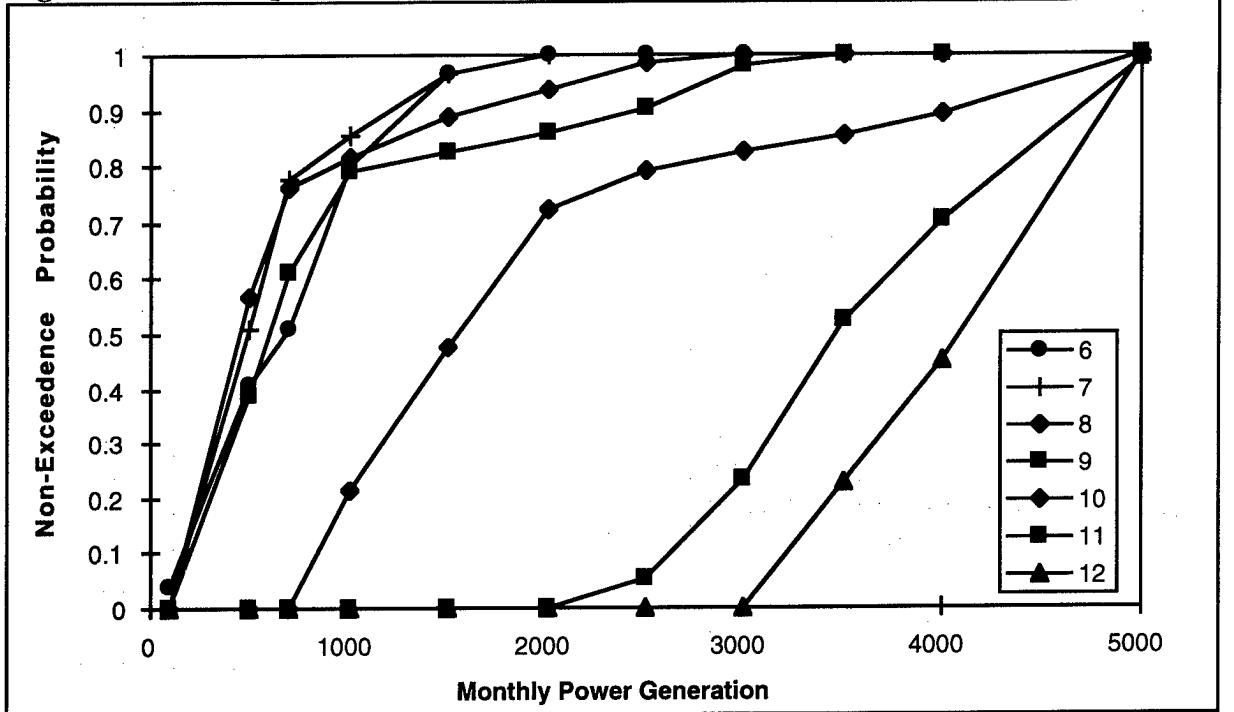


Figure 4.10: Example Non-Exceedence Probability Plots for Monthly Power Yields for June-December

4.9 Conclusions

The use of deterministic optimization models for seasonal reservoir operation planning begins with careful consideration of the appropriate seasonal operation storage, release and flow penalties, initial reservoir storage conditions, end-of-period storage targets and penalties, and relevant and available inflow hydrologies. For short seasonal or annual operation runs, all of these factors can be expected to significantly affect model results.

The use of optimization model results for seasonal operations must be "fit" into the institutional and engineering context of the reservoir operation problem. The use of optimization modeling must usually be a companion to current seasonal operation planning procedures, not a replacement for existing procedures. Optimization modeling must be used with the context of existing and available hydrologic data, streamflow forecasts, simulation models, water and power contracting periods, system operating rules, system operating purposes, and basin coordination agreements.

Several approaches are available for the use of optimization modeling to supplement existing simulation studies of seasonal operations. These include the use of optimization modeling results to incrementally modify longer-term operating rules for seasonal application, to prioritize drawdown and refill in multi-reservoir systems, to provide direct suggestions for short-term operations, and to speed provision of updated operating advice throughout the year. In addition, for broader reservoir management purposes, optimization model results can quickly provide estimates of water and power yield for each month during the coming year as well as storage and flow probability curves for recreation and navigation management activities.

Chapter 5

Conclusions

Several conclusions are suggested by this work:

1. *Reservoir systems typically require modeling studies to aid operational decision-making at several different time-scales.* Most typically, these include long-term strategic operations, seasonal operations, and real-time operations. Long-term strategic studies establish general operating approaches for the system and are conducted infrequently, every 5 to 10 years. Seasonal studies guide operations for the coming year or season, and are performed as frequently as several times a year. Real time operation studies establish very near term releases, on a daily or hourly basis, and can be conducted daily, weekly, or during particular episodes, such as floods. The wide range of potential operating strategies available for operating a reservoir system, as reviewed in Chapter 2, creates a need for standard and rigorous methods for developing and comparing alternative reservoir operation policies. Simulation and optimization modeling provide such methodologies, and are both highly cost-effective and rigorous if done properly.
2. *Deterministic optimization models can help establish operating policies during long-term, seasonal, and real-time operation time frames.* This report examined application formulation and results interpretation issues pertaining to the applying deterministic optimization models, such as HEC-PRM, to the establishment of long-term and seasonal operating rules. In general, these issues are tractable, and there are practical examples where these optimization models can provide useful results for developing seasonal and long-term operating policies, refining existing operating rules, and justifying the desirability of a particular set of operating rules. Optimization results also can help update operating policies within an operating season (USACE, 1996) or revise long-term operating policies on a strategic basis (USACE, 1992, 1994a, 1995).
3. *A variety of approaches are available for inferring strategic operating rules from long-term optimization results (such as HEC-PRM).* Chapter 3 examined over a half dozen approaches available for developing long-term or strategic reservoir operating policies from deterministic optimization results. Each of these approaches has been found to be useful for different contexts. However, all approaches generally require the use of simulation modeling for refinement and testing of operating rules.
4. *A wide variety of operating rule forms are available for making seasonal and long-term operation policies.* These rule forms are frequently supported by a substantial body of practical and theoretical experience. Existing operating rule forms, reviewed extensively in Chapter 2, apply to a wide range of reservoir system configurations and purposes, and can provide considerable flexibility for system operations.
5. *Optimization models such as HEC-PRM also can aid in seasonal operating decisions.* As a complement to simulation models for seasonal operations, optimization models can provide a rigorous second opinion suggesting promising modifications to seasonal operating rules for particular near-term conditions. These results can also provide probabilistic yield and performance measures for a number of reservoir management problems. Seasonal operation models also can be used several times a year to update operating plans as conditions and forecasts change (USACE, 1996).

6. *Simulation modeling is an essential companion to optimization modeling.* The shortcomings of direct use of optimization model results can generally be overcome by the use of simulation modeling for refining and testing the rules inferred from optimization model results. The advantages of simulation modeling include an ability to better represent the details of a system

and its objectives, realistic amounts of operator foresight into future inflows, and the more general acceptability of simulation model results. The optimization model adds a rigorous pursuit of the least-cost or maximum net benefit operating policy without prejudice from prior history or operations.

Appendix A References

Barrows, H. K. (1948), *Floods, Their Hydrology and Control*, McGraw-Hill Book Company, Inc., pp. 186, 280.

Bayazit, M. and N.E. Unal (1990), "Effects of Hedging on Reservoir Performance," *Water Resources Research*, Vol. 26, No. 1, April, pp. 713-719.

Bhaskar, N.R. and E.E. Whitlatch (1980), "Derivation of Monthly Reservoir Release Policies," *Water Resources Research*, Vol. 16, No. 6, December, pp. 987-993.

Bower, B.T., M.M. Hufschmidt, and W.W. Reedy (1966), "Operating Procedures: Their Role in the Design of Water-Resource Systems by Simulation Analyses," in A. Maass, et al., *Design of Water-Resource Systems*, Harvard University Press, Cambridge, MA, pp. 443-458.

Buras, N. (1963), "Conjunctive Operations of Dams and Aquifers," *Journal of the Hydraulics Division*, ASCE, Vol. 89, No. HY6, pp. 111-131.

Chung, F.I., M.C. Archer, and J.J. DeVries (1989), "Network Flow Algorithm Applied to California Aqueduct Simulation," *Journal of Water Resources Planning and Management*, ASCE, Vol. 115, No. 2, March, pp. 131-147.

Clark, E. J. (1950), "New York control curves," *Journal of American Water Works Association*, Vol. 42, No. 9, pp. 823-827.

Clark, E.J. (1956), "Impounding Reservoirs," *Journal of the American Water Works Association*, Vol. 48, No. 4, April, pp. 349-354.

Eltahir, E.A.B. (1996), "El Niño and the Natural Variability in the Flow of the Nile River," *Water Resources Research*, Vol. 32, No. 1, January, pp. 131-137.

Evenson, D.E. and J.C. Moseley (1970), "Simulation/Optimization Techniques for Multi-Basin Water Resource Planning," *Water Resources Bulletin*, Vol. 6, No. 5, September-October, pp. 725-736.

Ford, D.T., R. Garland, and C. Sullivan (1981), "Operation Policy Analysis: Sam Rayburn Reservoir," *Journal of the Water Resources Planning and Management Division*, ASCE, Vol. 107, No. WR2, October, pp. 339-350.

Ghosh, S. N (1986), *Flood Control and Drainage Engineering*, Oxford & IBH Publishing Co., pp. 158 - 159.

Hall, W.A. and J.A. Dracup (1970), *Water Resources Systems Engineering*, McGraw-Hill Publishing Co., New York, NY, 372 pp.

Hashimoto T., Stedinger J.R., and Loucks D.P. (1982), "Reliability, Resiliency, and Vulnerability Criteria for Water Resource System Performance Evaluation," *Water Resources Research*, Vol. 18, No. 1, February, pp. 14-20.

Hirsch, R.M. (1978), "Risk Analyses for a Water-Supply System-Occoquan Reservoir, Fairfax and Prince William Counties, Virginia", U.S. Geological Survey, Open-File Report 78-452.

Hirsch, R.M. (1981), "Estimating Probabilities of Reservoir Storage for the Upper Delaware River Basin", U.S. Geological Survey, Open-File Report 81-478.

Hirsch, R.M. (1981), "Stochastic Hydrologic Model for Drought Management," *Journal of the Water Resources Planning and Management Division*, ASCE, Vol. 107, No. WR2, October, pp. 303-314.

Hufschmidt, M.M. and M.B. Fiering (1966), *Simulation Techniques for Design of Water-Resource Systems*, Harvard University Press, Cambridge, MA.

Jacoby, H.D. and D.P. Loucks (1972), "Combined Use of Optimization and Simulation Models in River Basin Planning," *Water Resources Research*, Vol. 8, No. 6, December, pp. 1401-1414.

Jensen, M.R. and J.R. Lund (1993), "WATSUP: A Simulation Model for Water Supply Costs and Reliability," project completion report to the California Department of Water Resources, Sacramento, CA, June.

Jettmar, R.U. and G.K. Young (1975), "Hydrologic Estimation and Economic Regret," *Water Resources Research*, Vol. 11, No. 5, October, pp. 648-656.

Johnson, S.A., J.R. Stedinger, and K. Staschus (1991), "Heuristic Operating Policies for Reservoir System Simulation," *Water Resources Research*, Vol. 27, No. 6, May, pp. 673-685.

Karamouz, M. and M.H. Houck (1982), "Annual and Monthly Reservoir Operating Rules Generated by Deterministic Optimization," *Water Resources Research*, Vol. 18, No. 5, October, pp. 1337-1344.

Karamouz, M. and M.H. Houck (1987), "Comparison of Stochastic and Deterministic Dynamic Programming for Reservoir Operating Rules Generation," *Water Resources Bulletin*, Vol. 23, No. 1, February, pp. 1-9.

Karamouz, M., M.H. Houck, and J.W. Delleur (1992), "Optimization and Simulation of Multiple Reservoir Systems," *Journal of Water Resources Planning and Management*, ASCE, Vol. 118, No. 1, January/February, pp. 71-81.

Kelley, K.F. (1986), *Reservoir Operation During Drought: Case Studies*, Research Document No. 25, Hydrologic Engineering Center, Davis, CA.

Kelman, J., J.M. Damazio, J.L. Mariën, and J.P. daCosta (1989), "The Determination of Flood Control Volumes in a Multireservoir System", *Water Resources Research*, Vol. 25, No. 3, March, pp. 337-344.

King, I.P. and D.E. Evenson (1972), *The Development of Operating Rules for Multi-Basin Reservoir Systems*, Water Resources Engineers, Inc., Walnut Creek, CA, March.

Kirby, K.W. (1994), *Resolving Conflict Over Reservoir Operation: A Role for Optimization and Simulation Modeling*, M.S. Thesis, Department of Civil and Environmental Engineering, University of California, Davis, CA.

Klemes, V. (1974), "The Hurst Phenomenon: A Puzzle?", *Water Resources Research*, Vol. 10, No. 4, August, pp. 675-688.

Klemes, V. (1977), "Value of information in reservoir optimization," *Water Resources Research*, Vol. 13, No. 5, pp. 837-850.

Klemes, V. (1979), "Storage Mass-Curve Analysis in a Systems-analytic Perspective," *Water Resources Research*, Vol. 15, No. 2, April, pp. 359-370.

Kuszera, G. and G. Diment (1988), "General Water Supply System Simulation Model: WASP," *Journal of Water Resources Planning and Management*, ASCE, Vol. 114, No. 4, July, pp. 365-382.

Linsley, R.K., M.A. Kohler, and J.L.H. Paulus (1982), *Hydrology for Engineers*, McGraw-Hill Book Co., N.Y.

Loucks, D.P., P.N. French, and M.R. Taylor (1995), *IRAS - Interactive River-Aquifer System Simulation, Program Description and Operation*, School of Civil and Environmental Engineering, Cornell University, Ithaca, NY.

Loucks, D.P. and K.A. Salewicz (1989), *IRIS, an Interactive River System Simulation Program, General Introduction and Description*, School of Civil and Environmental Engineering, Cornell University, Ithaca, NY.

Loucks, D.P. and O.T. Sigvaldason (1982), "Multiple-reservoir operation in North America," in *The Operation of Multiple Reservoir Systems*, Z. Kaczmarek and J. Kindler (eds.), International Institute for Applied Systems Analysis, Laxenburg, Austria.

Loucks, D.P., J.R. Stedinger, and D.A. Haith (1981), *Water Resources Systems Planning and Analysis*, Prentice Hall, Inc., Englewood Cliffs, NJ.

Lund, J.R.(1995), "Derived Estimation of Willingness-to-Pay to Avoid Probabilistic Shortage," *Water Resources Research*, Vol. 31, No. 5, pp. 1367-1372, May.

Lund, J.R. and I. Ferreira (1996), "Operating Rule Optimization for the Missouri River Reservoir System," *Journal of Water Resources Planning and Management*, ASCE, Vol. 122, No. 4, July.

Lund, J.R., M. Jenkins, and O. Wilchfort (1995), "Integrated Planning and Management for Urban Water Supplies Considering Multiple Uncertainties," Project Completion Report for Project UCAL-WRC-W-813, University of California Water Resources Center, Davis, CA, September.

Maass, A., M.M. Hufschmidt, R. Dorfman, H.A. Thomas, Jr., S.A. Marglin, and G.M. Fair (1966), *Design of Water Resource Systems*, Harvard University Press, Cambridge, MA.

Mariën, J.L., J.M. Damázio, and F.S. Costa (1994), "Building flood control rule curves for multipurpose multireservoir systems using controlability conditions," *Water Resources Research*, Vol. 30, No. 4, April, pp. 1135-1144.

Martin, Q.W. (1987), "Optimal Daily Operations of Surface-Water Systems," *Journal of Water Resources Planning and Management*, ASCE, Vol. 113, No. 4, July, pp. 453-470.

Mejia, J.M., P. Egli, and André Leclerc (1974), "Evaluating Multireservoir Operating Rules," *Water Resources Research*, Vol. 10, No. 8, December, pp. 1090-1098.

Miquel, J. and P.A. Roche (1983), "Can an optimal release policy for a reservoir be obtained in the case of imperfect low flow forecasts?", (in French) in E. Plate and N. Buras (eds.), *Scientific Procedures Applied to the Planning, Design, and Management of Water Resources Systems*, IAHS Publ. no. 147, August, 1983.

Miquel, J. and P.A. Roche (1986), "La gestion des barrages réservoirs," *La Houille Blanche*, No. 6, pp. 409-425.

Moy, W.-S., J. L. Cohon , and C. S. ReVelle (1986), "A programming model for analysis of the reliability, resilience, and vulnerability of a water supply reservoir, *Water Resources Research*, Vol. 22, No. 4, pp. 489-498.

Rippl, W. (1883), "The Capacity of Storage-Reservoirs for Water Supply," *Minutes Proc. Inst. Civil Eng.*, Vol. 71, pp. 270-278.

Rogers, P.P. and M.B. Fiering (1986), "Use of Systems Analysis in Water Management," *Water Resources Research*, Vol. 22, No. 9, August, pp. 146S-158S.

Saad, M., P. Bigras, A. Turgeon, and R. Duquette (1996), "Fuzzy Learning Decomposition for the Scheduling of Hydroelectric Power Systems," *Water Resources Research*, Vol. 32, No. 1, January, pp. 179-186.

Saad, M. and A. Turgeon (1988), "Application of Principal Component Analysis to Long-Term Reservoir Management," *Water Resources Research*, Vol. 24, No. 7, pp. 907-912.

Saad, M., A. Turgeon, and J.R. Stedinger (1992), "Censored-data correlation and principal component dynamic programming," *Water Resources Research*, Vol. 28, No. 8, pp. 2135-2140.

Saad, M., A. Turgeon, P. Bigras, and R. Duquette (1994), "Learning Disaggregation Technique for the Operation of Long-Term Hydroelectric Power Systems," *Water Resources Research*, Vol. 30, No. 11, November, pp. 3195-3203.

Salas, J.D. and W.A. Hall (1983), "Operating Rules for Storage Reservoirs," in E. Guggino, G. Rossi, and D. Hendricks (eds.), *Operation of Complex Water Systems*, Martinus Nijhoff Publishers, Boston, MA, pp. 208-225.

Sand, G.M. (1984), *An Analytical Investigation of Operating Policies for Water-Supply Reservoirs in Parallel*, Ph.D. Dissertation, Cornell University, Ithaca, NY.

Sheer, D. (1986), "Dan Sheer's Reservoir Operating Guidelines," as interpreted by J.R. Stedinger, School of Civil and Environmental Engineering, Cornell University, June/July, 1 page.

Shih, J-S., and C. ReVelle (1994), "Water Supply Operations During Drought: Continuous Hedging Rule," *Journal of Water Resources Planning and Management*, Vol. 120, No. 5, September-October, pp. 613-629.

Shin, J-S. and C. Revelle (1995), "Water Supply operations during drought: A discrete hedging rule," *European Journal of Operational Research*, Vol. 82, pp. 163-175.

Sigvaldason, O.T. (1976), A simulation model for operating a multipurpose multireservoir system," *Water Resources Research*, Vol. 12, No. 2, pp. 263-278.

Stedinger, J. R. (1978), "Comment on 'Value of information in reservoir optimization,'" *Water Resources Research*, Vol. 14, No. 5, pp. 984-986.

Stedinger, J. R., B. F. Sule, and D. Pei (1983), "Multiple reservoir system screening models," *Water Resources Research*, Vol. 19, No. 6, pp. 1383-1393.

Stedinger, J. R., B. F. Sule, and D.P. Loucks (1984), "Stochastic Dynamic Programming Models for Reservoir Operation Optimization," *Water Resources Research*, Vol. 20, No. 11, pp. 1499-1505.

Su, M.-D., R. Narayanan, T.C. Hughes, and A.B. Bishop (1991), "Evaluation of Municipal Water Supply Operating Rules Using Stochastic Dominance," *Water Resources Research*, Vol. 27, No. 7, July, pp. 1519-1527.

Tejada-Guibert, J.A., J.R. Stedinger, and K. Staschus (1990), "Optimization of Value of CVP's Hydropower Production," *Journal of Water Resources Planning and Management*, ASCE, Vol. 116, No. 1, January, pp. 52-70.

Tejada-Guibert, J.A., S.A. Johnson, and J.R. Stedinger (1993), "Comparison of Two Approaches for Implementing Multi-Reservoir Operating Policies Derived Using Dynamic Programming," *Water Resources Research*, Vol. 29, No. 12, pp. 3969-3980.

Tejada-Guibert, J.A., S.A. Johnson, and J.R. Stedinger (1995), "The Value of Hydrologic Information in Stochastic Dynamic Programming Models of a Multireservoir System," *Water Resources Research*, Vol. 31, No. 10, pp. 2571-2579.

Toebe, G.H. and C. Rukvichai (1978), "Reservoir System Operating Policy - Case Study," *Journal of the Water Resources Planning and Management Division*, ASCE, Vol. 104, No. WR1, November, pp. 175-191.

Trott, W.J. (1979), *Determination of Real-Time Operating Rules for a Multiple Reservoir System*, Ph.D. dissertation, University of California, Los Angeles, CA.

Turgeon, A. (1981), "A Decomposition Method for the Long-Term Scheduling of Reservoirs in Series," *Water Resources Research*, Vol. 17, No. 6, December, pp. 1565-1570.

U.S. Army Corps of Engineers (1976), *Flood Control by Reservoirs*, Hydrologic Engineering Methods for Water Resources Development, Volume 7, Hydrologic Engineering Center, Davis, CA, February.

U.S. Army Corps of Engineers (1977), *Reservoir System Analysis for Conservation*, Hydrologic Engineering Methods for Water Resources Development, Volume 9, Hydrologic Engineering Center, Davis, CA, June.

U.S. Army Corps of Engineers (1979a), *Missouri River Main Stem Reservoirs System Reservoir Regulation Manual: Volume 1, Master Manual*, Missouri River Division, Omaha, Nebraska.

U.S. Army Corps of Engineers (1979b), *Flood Tolerance in Plants: A State-of-the Art Review*, Technical Report E-79-2, Waterways Experiment Station, Vicksburg, MS.

U.S. Army Corps of Engineers (1985), *Engineering Manual: Engineering and Design, Hydropower*, EM 1110-2-1701, U.S. Army corps of Engineers, Washington D.C.

U.S. Army Corps of Engineers (1987), *Management of Water Control Systems*, EM 1110-2-3600, U.S. Army Corps of Engineers, Washington, D.C., 30 November.

U.S. Army Corps of Engineers (1991a), *Missouri River System Analysis Model - Phase I*, Report PR-15, Hydrologic Engineering Center, Davis, CA, February.

U.S. Army Corps of Engineers (1991b), *Economic Value Functions for Missouri River System Analysis Model - Phase I*, Institute For Water Resources, Fort Belvoir, VA.

U.S. Army Corps of Engineers (1991c), *Authorized and Operating Purposes of Corps of Engineers Reservoirs*, preliminary report, U.S. Army Corps of Engineers, Hydrologic Engineering Center, Davis, CA.

U.S. Army Corps of Engineers (1991d), *Optimization of Multi Purpose Reservoir System Operations: A Review of Modeling and Analysis Approaches*, Research Document No. 34, U.S. Army Corps of Engineers, Hydrologic Engineering Center, Davis, CA, January.

U.S. Army Corps of Engineers (1992), *Developing Operation Plans from HEC Prescriptive Reservoir Results for the Missouri River System: Preliminary Results*, Report PR-18, Hydrologic Engineering Center, Davis, CA, March.

U.S. Army Corps of Engineers (1993a), *Columbia River Reservoir System Analysis: Phase II*, Report PR-21, Hydrologic Engineering Center, Davis, CA, December.

U.S. Army Corps of Engineers (1993b), *Missouri River Master Water Control Manual: Preliminary Draft*, Missouri River Division, Omaha, Nebraska.

U.S. Army Corps of Engineers (1994a), *HECPRM: Hydrologic Engineering Center's Prescriptive Reservoir Model, Program Description*, Hydrologic Engineering Center, Davis, CA, February.

U.S. Army Corps of Engineers (1994b), *Operating Rules from HEC Prescriptive Reservoir Model Results for the Missouri River System: Development and Preliminary Testing*, Report PR-22, Hydrologic Engineering Center, Davis, CA, May.

U.S. Army Corps of Engineers (1994c), *Hydrologic Engineering Center Project Benefit Accomplishment Package (HEC-PBA)*, Report CPD-67, Hydrologic Engineering Center, Davis, CA, September.

U.S. Army Corps of Engineers (1995), *Preliminary Operating Rules for the Columbia River System from HEC-PRM Results*, Report PR-26, Hydrologic Engineering Center, Davis, CA, June.

U.S. Army Corps of Engineers (1996), *Application of HEC-PRM for Seasonal Reservoir Operation of the Columbia River System*, Hydrologic Engineering Center, Davis, CA, June.

Varlet, M. (1923), "Étude Graphique des Conditions d'Exploitation d'un Réservoir de Régularisation," *Annales des Ponts et Chaussées*, Partie Technique, 93e Année, 11e Série-Time 64, Fasc. 4. Tome 2, July-August.

Viessman, W., J.K. Knapp, G.L. Lewis, and T.E. Harbaugh (1977), *Introduction to Hydrology*, Harper and Row, Publishers, N.Y.

Whitlatch, E.E. and N.R. Bhaskar (1978), *Application of Mathematical Optimization Techniques in Reservoir Design and Management Studies*, Project Completion Report No. 525X, Water Resources Center, Ohio State University, Columbus, OH, December.

WRMI (1994), "The WRMI Modeling System: Meeting Future Water Needs," *WRMILink*, Water Resources Management, Inc., Columbia, MD, Vol. 3, No. 16, January, p. 2.

Yeh, W. W-G. (1982), *State of the Art Review: Theories and Applications of Systems Analysis Techniques to the Optimal Management and Operation of a Reservoir System*, Technical Report No. UCLA-ENG-82-52, UCLA School of Engineering and Applied Science, Los Angeles, CA, June.

Young, G.K. (1966), *Techniques for Finding Operating Rules*, Ph.D. Thesis, Harvard University, Cambridge, MA, 100 pp; a derivative work was published as "Finding Reservoir Operating Rules," *Journal of the Hydraulics Division*, ASCE, Vol. 93, No. HY6, November, pp. 297-319.

Appendix B

Summary of Columbia River System

Seasonal Operation Study¹

Report Summary

This report presents the results and conclusions of an application of the Hydrologic Engineering Center's Prescriptive Reservoir Model (HEC-PRM) for seasonal operation of the Columbia River System. A position analysis approach is used to suggest promising seasonal operations for the Columbia River System which can be updated throughout the annual drawdown refill cycle. Such HEC-PRM-based seasonal reservoir operation advice could offer guidance in simulation testing and reduce the number of simulation runs needed to formulate seasonal operation plans.

HEC-PRM is run using the position analysis approach, a common form of risk analysis designed to examine reservoir operations for seasonal periods (Hirsch, 1978). Position analysis addresses seasonal operation rather than long-term, strategic operation. The procedure uses a simulation or optimization model to conduct separate runs for many (n) scenarios of future seasonal hydrologies. Each model run begins with the same, current reservoir storage. The number of runs (n) is determined by the number of inflow sequences available, based usually on n years of historical record or n alternative forecasts for future inflows.

Although greatly modified in recent years, due to environmental concerns, the Columbia River System traditionally operates on a seasonal basis. The three operating seasons include the fixed drawdown season (August-December), variable drawdown season (January-March), and refill season (April-July). Each year, hundreds of simulation model (HYSSR) runs are conducted to plan seasonal operations. Four HEC-PRM seasonal studies are presented in this report. Each study captures at least one of these three traditional operating seasons.

This project is the first extensive use of HEC-PRM as a seasonal reservoir operation model. Past HEC-PRM studies of the Columbia River System are strategic planning studies (USACE, 1991, 1993, 1995). The idea of using HEC-PRM as a seasonal model was proposed and encouraged by a preliminary HEC-PRM seasonal study in 1995 (USACE, 1995).

The findings of this report demonstrate that HEC-PRM is potentially useful for seasonal operation studies of the Columbia River System. Overall, for the four studies in this report, the HEC-PRM seasonal operation advice is reasonable and consistent. Simulation should be used to refine and test HEC-PRM seasonal advice to explore its potential for improved operations. The use of HEC-PRM may allow for a considerable focusing of detailed simulation studies.

¹From the U.S. Army Corps of Engineers (1996), *Application of HEC-PRM for Seasonal Reservoir Operation of the Columbia River System*, "Executive Summary," HEC, Davis, CA, June.

Background

Columbia River System

The Columbia River System is located in the Pacific Northwest region of the United States (Figure B.1). The entire Columbia River System is comprised of over 250 reservoirs and 100 hydroelectric projects. For the HEC-PRM seasonal operation studies in this report, the reservoir system is represented by a selection of key reservoirs only. Figure B.2 shows the network developed for HEC-PRM runs. This reservoir network was formulated in previous USACE Columbia River reports (USACE, 1991, 1993, 1995).

Seven main storage reservoirs are the focus of the seasonal operation study analysis. The seasonal reservoir operations for Mica, Arrow, Grand Coulee, Duncan, Libby, Hungry Horse and Dworshak reservoirs are the operations discussed throughout the four studies in this report.

Inflow Hydrology

Standardized inflow hydrology for the period of 1928 - 1978 is used in each seasonal study (USACE, 1993). Low and high flow patterns are present and critical periods are included. The standardized inflows are adjusted to reflect 1980 depletions and are modified to incorporate inflow forecasts when available.

The forecast modifications made to historical inflows were performed by the U. S. Army Corps of Engineers North Pacific Division (USACE NPD). Inflow forecasts are made at the beginning of the month, for the months of January to June. As a result, the inflow hydrology for the 1994 Drawdown season study, which spans from July to March, is not modified because inflow forecasts are unavailable.

Approach Overview

The approach to seasonal operation studies presented here uses HEC-PRM according to the position analysis technique.

“Position analysis is a specialized application of risk analysis. Its purpose is to estimate the risks associated with a given plan of operation over a period of a few months...it consists of n separate simulations rather than one continuous simulation of length n years. Each of these simulations is initialized with the same reservoir storage value--that storage actually existing in the reservoir at the beginning of the present month. Thus, it is an analysis of risks evaluated from the present 'position'" (Hirsch, 1978).

Hirsch discusses the use of the position analysis approach for simulation modeling only. Position analysis also can be applied to optimization studies, evidenced by the HEC-PRM seasonal studies using position analysis in this report. The use of position analysis with an optimization model allows for rapid identification of promising short-term operating advice for consideration by system operators and more detailed simulation testing and refinement.

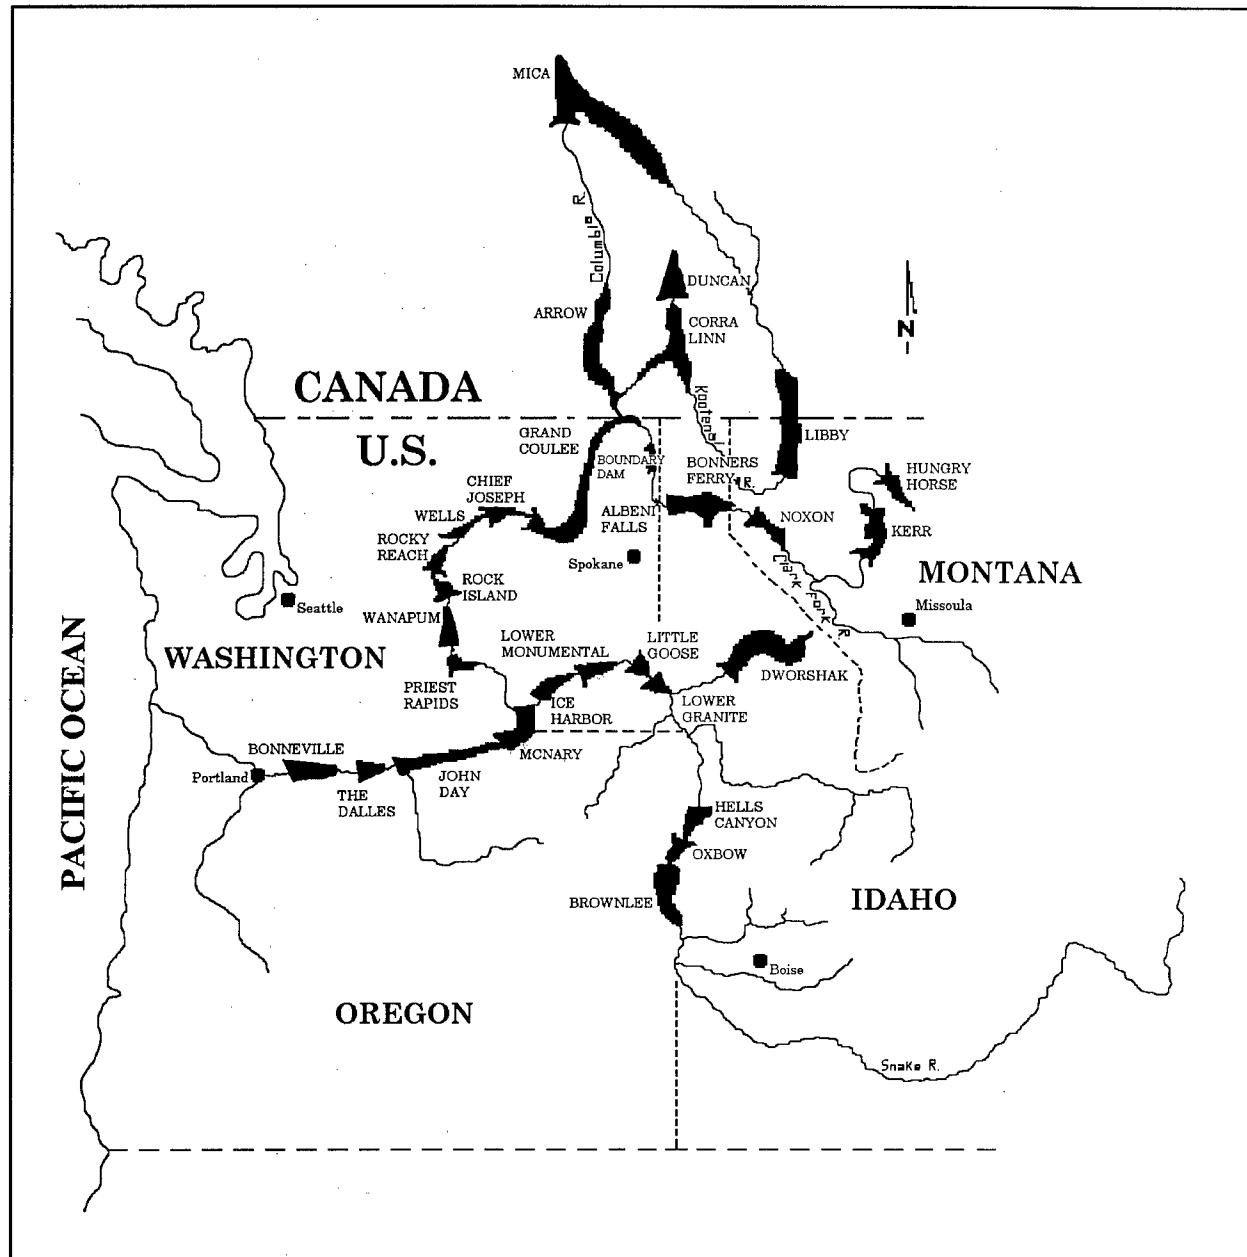


Figure B.1: Columbia River System

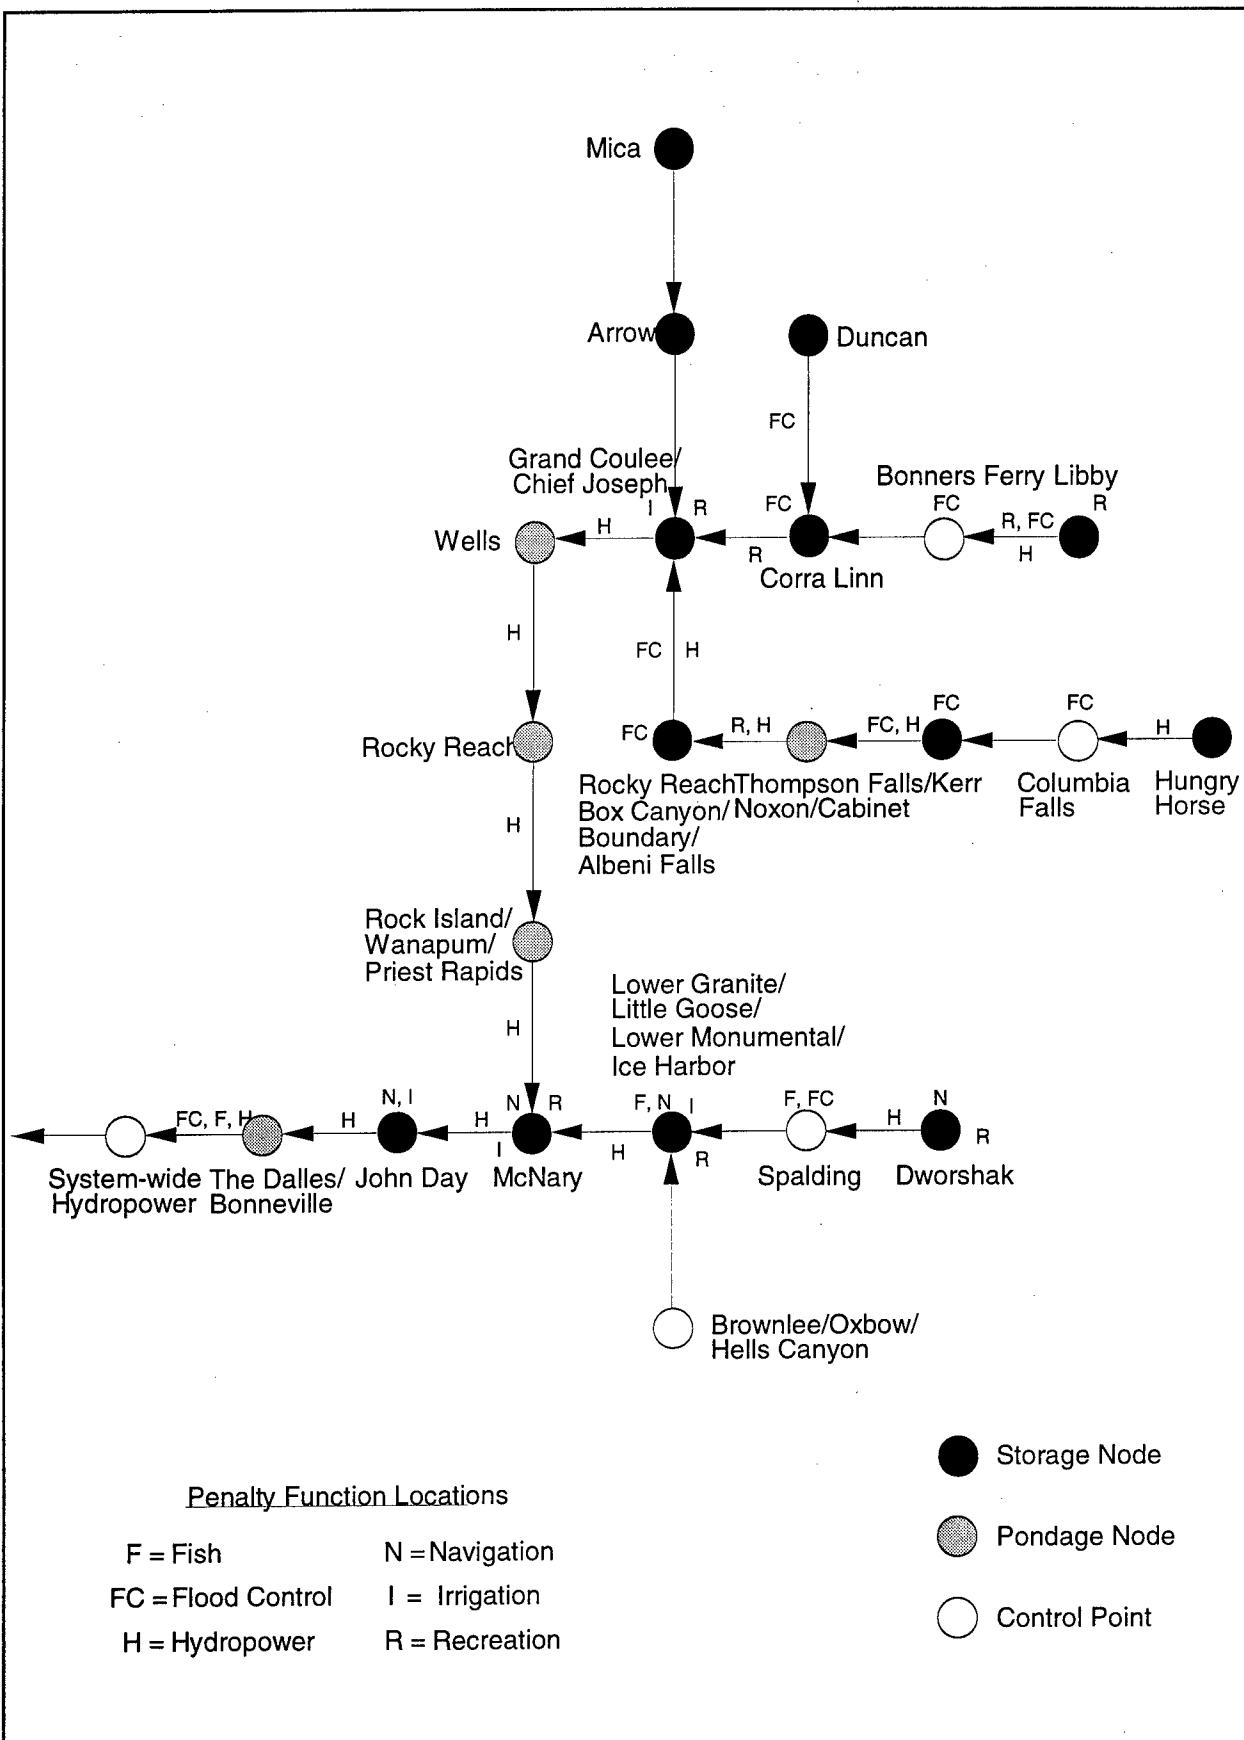


Figure B.2: Updated Single-Period Network Model of Columbia River System

HEC-PRM is a network flow model that optimizes reservoir operations by minimizing flow and storage penalties or costs throughout a reservoir system network. Position analysis, as stated above, is a common study approach that focuses on short-term, seasonal periods, conducting many separate model runs for a range of historically-based future flow scenarios. In its most basic form, n years of historical record are divided to provide data for n shorter runs of a seasonal operations model (Hirsch, 1978).

The number of runs (n) is directly related to the number of inflow sequences available. In this report, the historical streamflow record forms the basis for at least 48 seasonal forecasts of system inflows. These inflow scenarios are then used in at least 48 separate HEC-PRM runs to find the ranges of promising operations for this system. As explained by Hirsch, for each model run, each reservoir begins at a given current initial storage, or "position."

HEC-PRM

HEC-PRM is the model used to suggest seasonal reservoir operations for the four seasonal studies in this report. HEC-PRM is a prescriptive (or optimization) model and, therefore, the model optimizes the allocation of available water in the Columbia River System to find seasonal reservoir operations. HEC-PRM also is a network flow model. As a result, a network of nodes (reservoirs) and links (channels, diversions, etc.) needs to be defined to represent the actual, physical framework of a reservoir system, the Columbia River System in this case (Jensen and Barnes, 1980).

As a prescriptive model, HEC-PRM finds solutions based on predetermined operational objectives. Penalty functions define these operational objectives. The objective function of the network flow problem is the sum of the convex, piecewise-linear approximations of the penalty functions (USACE, 1991).

The aim of the use of HEC-PRM is to develop storage and release advice for use in more detailed simulation studies and to decrease the number of simulation runs required to formulate seasonal reservoir operation plans. Advantages of using HEC-PRM for this purpose are that the quantity of runs necessary to reach a storage or release target is typically less than for a simulation model and the model is driven explicitly by formally stated system operating purposes, in the form of penalty functions.

A limitation of HEC-PRM is the model's omniscient perspective of future inflows. This allows HEC-PRM perfect foresight into future seasonal inflows, which is unrealistic, and, therefore, subsequent simulation testing is usually required.

Step-by-Step Seasonal Study Procedure

The step-by-step seasonal study procedure using HEC-PRM with the position analysis approach for the Columbia River System is as follows.

1. *Develop a HEC-PRM model of the system.* This includes representing the actual reservoir system as a network of nodes and links. Penalty functions are formulated to drive the

optimization process and define the operating objectives of the system, both economic and non-economic. Both the reservoir network and penalty functions were already developed for the Columbia River System when these seasonal studies were begun (USACE, 1991, 1993, 1995).

2. *Define the operating seasons of the reservoir system.* For the Columbia River System, there are three operating seasons: the fixed drawdown season (August-December), the variable drawdown season (January-March), and the refill season (April-July).

3. *Define the seasonal periods for each optimization study.* More than one season may be included in a seasonal study. Four seasonal studies are presented in this report. Three of the four seasonal studies in this report, the 1994 and 1995 January - July studies and the 1994 Drawdown season study, span two of the three operating seasons in the Columbia River System. The 1994 and 1995 January - July studies incorporate both the variable drawdown season and the refill season. The 1994 Drawdown season study encompasses the fixed drawdown season and the following variable drawdown season. The 1995 April - July seasonal update study covers the refill season only.

4. *Formulate end-of-period storage penalty functions.* End-of-period storage penalty functions manage carryover storage at the end of each study period. For the Columbia River System end-of-period storage penalty functions, the median storage results from a USACE NPD simulation model (HYSSR) study in 1995 were used as storage targets, with penalties for missing this target equal to the value of stored energy (USACE, 1995).

5. *Set current initial storage values for each reservoir.* Here, Actual Energy Regulation (AER) storage values were used in the seasonal studies to represent the initial storages or starting "positions" of each reservoir.

6. *Specify the inflow hydrologies to be used in the seasonal operation study.* Historical flows and forecasted inflows are used throughout the four seasonal studies in this report. The forecasted inflows are historical inflows modified by flow forecasts. The flow forecasts are determined monthly according to snowpack and soil moisture conditions. These forecasts are only available from January to June in the Columbia River System. As a result, forecasted inflows were available for every study in this report except the 1994 Drawdown season study, where historical flows were used.

7. *Run HEC-PRM for each inflow sequence.* For the seasonal studies in this report, the number of years of inflow available for each study ranged from 47 to 50 years. Each reservoir starts at the current initial storage, or "position," and the optimization analysis is run for the length of the season of interest.

8. *Interpret the HEC-PRM storage and release results.* Numerous graphs are used to aid interpretation of results. Position analysis plots show the storage or release results for each run overlaid upon each other; this display clearly shows the band of storage and release results suggested by HEC-PRM given the initial storage and range of inflow hydrologies. Quartile plots are a statistical representation of the position analysis plots; only the minimum, maximum and 25th, 50th and 75th percentile storage or release results are plotted. Exceedance and non-

exceedance plots and storage allocation graphs also are used to evaluate the HEC-PRM seasonal reservoir operations. Storage allocation plots are useful to determine basic refill or drawdown operations on a system-wide basis. The intent is to examine the optimization results to find consistent and promising near-term advice for efficient operations.

9. *Test HEC-PRM advice from the study conclusions with simulation.* The HEC-PRM advice should be able to direct the focus of simulation studies and lessen the number of simulation runs required to establish seasonal reservoir operation plans. This part of a seasonal reservoir study was not conducted for the studies in this report.

Seasonal Operation Application with Many Flow Forecasts

This section discusses the HEC-PRM seasonal operations for the Columbia River System for seasons in which flow forecasts are available. Forecasted inflows are available only from January to June. Many flow forecasts are made each month during this period from current snowpack and moisture conditions, allowing for possible modifications to each year of the historical inflow record. Three of the four studies discussed in this report have forecasts for the seven reservoirs under study available for use. The 1994 and 1995 January 1 flow forecasts are used in the 1994 and 1995 January - July studies respectively. Similarly, the April 1 inflow forecasts are used in the 1995 April - July seasonal update study. Since the 1994 Drawdown season study begins in July, flow forecasts are unavailable and historical inflows are used for the optimization analysis.

The HEC-PRM results for each of the seasonal studies in this report are analyzed to provide the following six items.

1. The probability of refill or drawdown for each of the seven reservoir is examined because a main goal of optimization and simulation modeling is to suggest how to operate a reservoir system to reach the end-of-period target storage. Analysis is conducted to assess if each reservoir reaches its target storage at the end of the season for all inflow sequences.

2. The HEC-PRM system-wide operation of the reservoirs is compared to the Actual Energy Regulation (AER) operation of the reservoir system. AER storages were used for the initial reservoir storages. The HEC-PRM system reservoir operation should be fairly similar to the operation used for the initial reservoir storages to ensure that HEC-PRM produces realistic seasonal operations.

3. The HEC-PRM system-wide storage allocation is examined to discover HEC-PRM's advice on system-wide drawdown or refill. Storage allocation analysis shows the order of reservoir drawdown or refill desirable for seasonal operations.

4. HEC-PRM and AER storage trends (drawdown, refill or level storage) from month-to-month are compared. Storage trend comparisons show if HEC-PRM operates each reservoir with the same basic trend as the AER operation. For the 1994 Drawdown study, HYSSR storage trends are available for comparison also.

5. Study the storage magnitude difference between HEC-PRM storage values and the AER operation. It is important to know the variation between HEC-PRM storage operation and the established operation, such as AER, for a seasonal period.

6. HEC-PRM specific quantitative storage and release results are determined. Any strong HEC-PRM quantitative advice is potentially useful for input into simulation studies.

Near-Term Period Analysis

Seasonal operation study result analysis typically focuses on the near-term period within each study. The "near-term" here is the first three months in a seasonal study. For instance, in the 1995 January - July study, the majority of the result analysis focuses on the January - March period. Near-term analysis is emphasized because of the potential use of seasonal update studies, conducting new optimization studies every month or every several months, where only the near-term information is valuable. Seasonal update studies are seasonal operation studies re-run within a seasonal period as current storage conditions and inflow forecasts are updated. The 1995 April - July seasonal update study in this report explores this technique.

January - July Season Results (1994 and 1995)

The 1994 and 1995 January - July studies both examine the variable drawdown season and refill season, but with different inflow forecasts and slightly different starting storage conditions. As a result, a comparison can be made between HEC-PRM results for both studies, understanding that initial reservoir storages and forecasted inflow hydrology vary. There is a significant difference between the inflow characteristics of these two studies; the 1994 forecasted inflows are less than those in 1995. Actually, the 1994 water year was drier than the 1995 water year (CRWMG, 1994).

The six key items focused on in the analysis of the HEC-PRM results are compared below for the 1994 and 1995 January - July studies.

1. HEC-PRM refilled more system reservoirs to their target storage for all of the inflow sequences tested in the 1995 study than the 1994 study. Four reservoirs, Mica, Arrow, Grand Coulee and Libby, always stored the target level in July in the 1995 study. Only two reservoirs, Mica and Grand Coulee, always reached their target storages in the 1994 January - July study. These findings are logical since less water was forecasted in 1994.

2. For both January - July season studies, HEC-PRM suggests drawing down the system throughout the near-term period (January, February and March, coincidentally the variable drawdown season). Notably, HEC-PRM suggests storing less water at the end of March in 1995 than 1994; HEC-PRM is aware that the greater 1995 forecasted inflows will refill the system sufficiently.

The 1994 Actual Energy Regulation operations draw down the system much as HEC-PRM did. However, 1995 actual AER operations were not so consistent throughout January - March (Figure B.3), beginning refill in February.

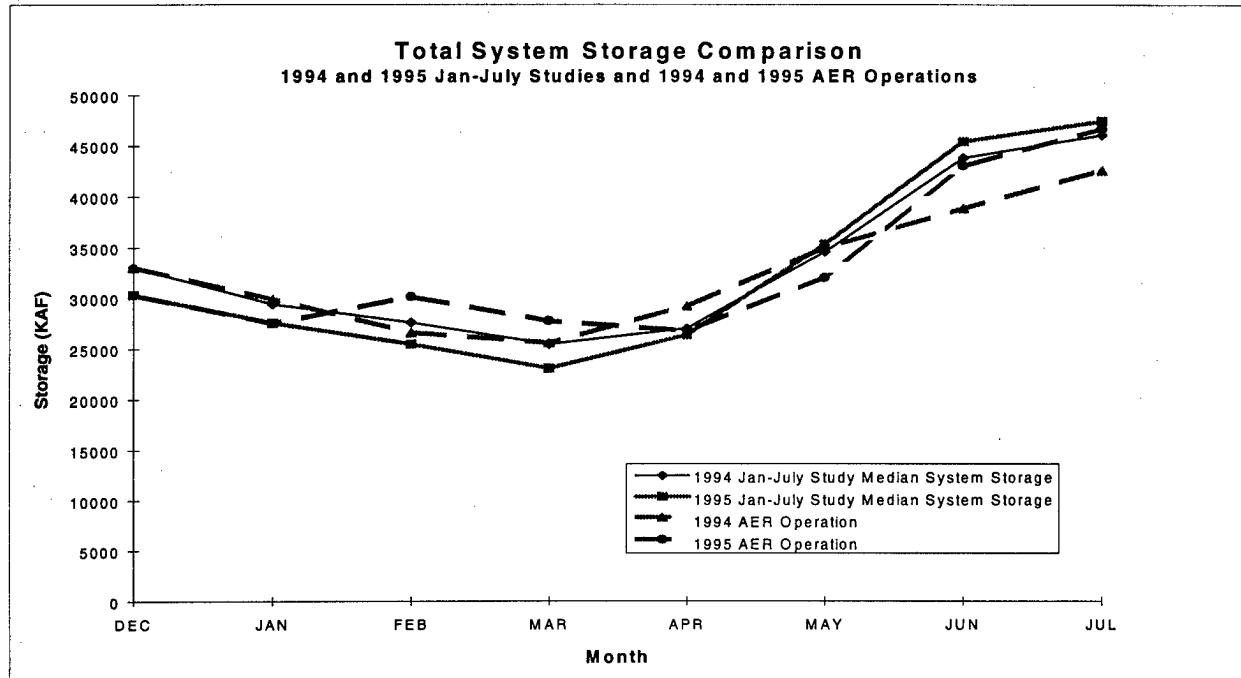


Figure B.3: Comparison of Total System Storage for 1994 and 1995 Jan-July Studies and 1994 and 1995 AER Operations

3. HEC-PRM suggests similar system-wide ordering of reservoir drawdowns (storage allocation) for the January - March variable drawdown operation for both 1994 and 1995 studies. Arrow reservoir is drawn down first, followed by either Mica or Dworshak. Grand Coulee is drawn down fourth. In the 1995 January - July study, Libby reservoir is the last reservoir to be drawn down, while both Libby and Hungry Horse draw down last together in the 1994 study.
4. The storage trends of HEC-PRM and AER compare better for the 1994 operations than the 1995 operations. For instance, 13 of 21 possible storage trends match for 1994, while 9 of 21 trends agree between the 1995 HEC-PRM and AER operations. In addition, comparing the HEC-PRM storage trends together, 16 of 21 trends agreed (Table B.1).
5. Comparison of the storage magnitude between HEC-PRM results and actual AER operations shows stronger agreement between them in the 1995 variable drawdown season than the 1994 variable drawdown season. Furthermore, given HEC-PRM's tendency to draw down the system more in 1995 (the wetter water year) than in 1994, HEC-PRM also tends to store less water in a number of reservoirs in 1995 than the 1995 AER operation. Conversely, in 1994, HEC-PRM tends to store more water than the AER operation in more reservoirs, responding to lesser forecasted inflows.
6. HEC-PRM specific quantitative storage and release advice is strong for both 1994 and 1995 studies. In both January - July season studies, HEC-PRM suggests releasing 603KAF, the minimum release, from Mica each month of the variable drawdown season. Both HEC-PRM studies store the minimum allowable storage of 227KAF in Arrow monthly from January to

Table B.1: Comparison of Storage Trends for 1994 and 1995 Variable Drawdown Seasons

RESERVOIR	1995 Jan - July Study	1994 Jan - July Study
January		
Mica	Drawdown	Drawdown
Arrow	Drawdown	Drawdown
Grand Coulee	Refill	Refill
Duncan	Refill	Refill
Libby	Drawdown	Drawdown
Hungry Horse	Refill	Drawdown
Dworshak	Drawdown	Drawdown
February		
Mica	Drawdown	Drawdown
Arrow	Maintain 227KAF	Drawdown
Grand Coulee	Drawdown	Variable
Duncan	Refill	Refill
Libby	Variable	Drawdown
Hungry Horse	Drawdown	Drawdown
Dworshak	Drawdown	Drawdown
March		
Mica	Drawdown	Drawdown
Arrow	Maintain 227KAF	Maintain 227KAF
Grand Coulee	Drawdown	Drawdown
Duncan	Refill	Refill
Libby	Drawdown	Drawdown
Hungry Horse	Variable	Drawdown
Dworshak	Drawdown	Drawdown

March. HEC-PRM suggests the following releases in January, February, and March in both January - July studies: Duncan reservoir at 6KAF (minimum allowable release) per month, Libby releases of 181KAF, the minimum allowable release, each month, Hungry Horse releases of 60KAF monthly and Dworshak releases between 300KAF and 450KAF each month. Additionally, for the 1994 January - July study, Grand Coulee stores 9107KAF (maximum storage).

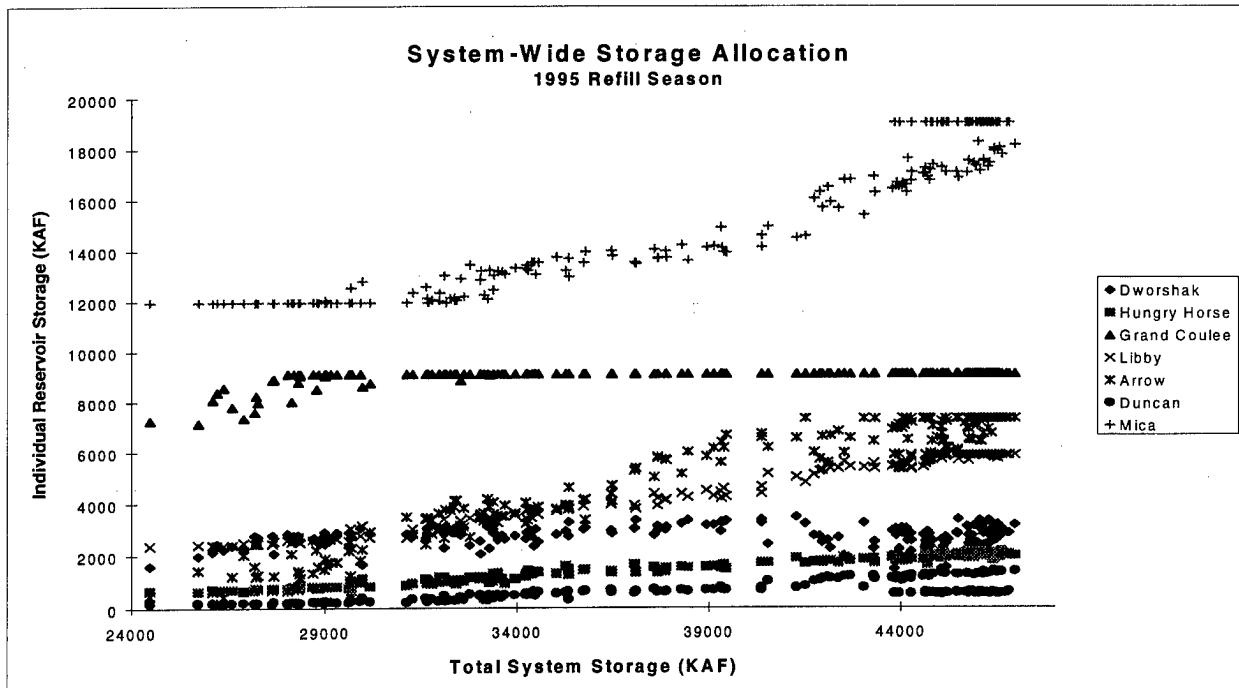
April - July Season Results (1995)

The 1995 April - July study is a seasonal update study for the 1995 January - July period. Updated inflow forecasts and storage levels of April 1st were used to run the 1995 April - July study for the 1995 refill study. Comparison of the HEC-PRM 1995 refill operations for the 1995 April - July seasonal update study and the 1995 January - July study shows that HEC-PRM refill operations were modified in the seasonal update study.

Specifically, the 1995 April - July study operations follow the AER operation more closely than the 1995 January - July study results. It is encouraging that HEC-PRM offers reservoir operation modifications as new flow and storage information becomes available. Therefore, HEC-PRM seasonal update studies are feasible for continuous improvement of seasonal operations, given new inflow forecasts and storage updates.

Below, the six main findings from the 1995 April - July seasonal update study results are presented.

1. HEC-PRM refilled three of the seven storage reservoirs, Mica, Grand Coulee and Libby reservoirs, to their target storages for all inflow sequences in the 1995 April - July study. The number of reservoirs that HEC-PRM always refilled to their target storages decreased by one from the 1995 January - July study to the 1995 April - July study. HEC-PRM always refilled Arrow reservoir in the 1995 January - July, but, with the updated inflow forecasts and storage levels of April, HEC-PRM clearly did not have enough water to ensure that Arrow reservoir always would reach its target storage in the 1995 April - July study.
2. The system-wide operations for both HEC-PRM and AER operations are refill in the 1995 April - July study. HEC-PRM's April to June system-wide operation in the 1995 January - July study is the same: consistent refill. Notably, HEC-PRM's system-wide storage is closer to the AER operation for the 1995 seasonal update study than the 1995 January - July study.
3. HEC-PRM allocates storage and orders refill among the seven reservoirs in the 1995 April - July as follows (Figure B.4). Grand Coulee reservoir refills first to the 9107KAF level. Arrow and Mica reservoirs significantly refill next. Duncan, Libby, Hungry Horse and Dworshak reservoirs are refilled after Grand Coulee, Arrow and Mica reservoirs begin refilling. The similarities between HEC-PRM's storage allocation for the 1995 January - July study and the 1995 April - July study are few. Mica and Arrow reservoirs refill once Grand Coulee refills to 9107KAF and levels off. A discrepancy between the two studies is that Grand Coulee is first priority for refill in the seasonal update study but Libby refills first in the 1995 January - July study.



**Figure B.4: System-Wide Storage Allocation for Refill for HEC-PRM 1995
Apr-July Study**

4. HEC-PRM and AER storage trends match for 12 of 21 comparisons for the 1995 seasonal update study.

5. HEC-PRM typically stores more water in Mica, Grand Coulee, Duncan, Libby and Hungry Horse reservoirs in April, May and June than the AER operation in the 1995 April - July study. Similarly, in the 1995 January - July study, HEC-PRM stores more water in the above five reservoirs than the AER operation.

6. HEC-PRM's specific quantitative storage and release advice is strong for the 1995 April - July seasonal update study (Table B.2). Grand Coulee should store 9107KAF in April, May and June. Arrow and Duncan releases of 302KAF and 6KAF, respectively, should be made all three months.

Libby and Hungry Horse releases for April, May and June are 181KAF and 60KAF, respectively. Dworshak releases should range from 300KAF to 450KAF each month. The specific quantitative advice from HEC-PRM is the same between the 1995 April - July study and the 1995 January - July study for Grand Coulee, Duncan, Libby, Hungry Horse and Dworshak for all three months.

Table B.2: Comparison of HEC-PRM Specific Advice (KAF) for Both 1995 Studies

Mica	1995 Apr-July	%	1995 Jan-July	%
April	Store 11950	50	Store 14075	50
May	Release 0	50	Release 0 - 145	25
June	Release 0	75	Release 0 - 150	50
Arrow				
April	Release 302(Min)-771	25	Store 227	50
May	Release 302	75	SAME	50
June	Release 302	25	SAME	50
Grand Coulee				
April	Store 9107(Max)	50	SAME	75
May	Store 9107	75	SAME	100
June	Store 9107	100	SAME	100
Duncan				
April	Release 6(Min)	100	SAME	100
May	Release 6	100	SAME	100
June	Release 6	100	SAME	100
Libby				
April	Release 181(Min)	100	SAME	25
May	Release 181	100	SAME	25
June	Release 181	75	SAME	25
Hungry Horse				
April	Release 60	75	SAME	75
May	Release 60	100	SAME	75
June	Release 60	75	SAME	75
Dworshak				
April	Release 300-450	50	SAME	75
May	Release 300-450	75	SAME	75
June	Release 300-450	75	SAME	50

Seasonal Operation Application without Flow Forecasts

The 1994 Drawdown study is unique because the fixed drawdown season is the only season in the Columbia River System without flow forecasts. Inflow forecasting from snowpack and soil moisture conditions is not available from July to December. As a result, the season from July to December is typically operated according to a fixed drawdown pattern and the 1994 Drawdown season study is run using historical inflow hydrology.

The 1994 Drawdown season study results are analyzed the same ways that the 1994 and 1995 January - July studies and the 1995 April - July study were analyzed. As a result, the six key items described earlier in the "Seasonal Operation Application with Many Flow Forecasts" section were the focus of the result analysis and they are presented below. Notably, the 1994 Drawdown season study results are compared to the AER operation as usual, and HYSSR simulations operations as well. The 1994 Drawdown season study is the only seasonal study in this report for which HYSSR results were available for comparison.

1. HEC-PRM always drew down all seven reservoirs to their respective March target storages.
2. Both HEC-PRM and actual AER operations begin system-wide drawdowns in August, while HYSSR starts system-wide drawdown in September. HEC-PRM typically stores more water in the system than the AER operation. HEC-PRM and HYSSR system-wide storages tend to overlap with a slight tendency for HEC-PRM to store a small amount more water than HYSSR.
3. HEC-PRM allocates storage by drawing down Mica, Duncan and Dworshak reservoirs first. Grand Coulee is kept high and level at 9107KAF as long as possible. Arrow reservoir is drained to 227KAF. Consequently, Grand Coulee is drawn down significantly. Mica, Libby, Hungry Horse and Dworshak reservoirs draw down also. Duncan stays relatively level after its initial drawdown to 30KAF (minimum allowable storage).
4. HEC-PRM and AER storage trend operations match for 11 of 21 comparisons. HEC-PRM and HYSSR storage trends agree for 12 of 21 comparisons. All three operations have similar storage trends for 7 of 21 instances.
5. HEC-PRM typically stores more water in Mica reservoir than HYSSR in the near-term, but approximately the same amount as the AER operation (Figure B.5). Among the three operations, HEC-PRM stores the least amount of water in Arrow in July, August and September. Grand Coulee is operated at or near 9107KAF by all three operations throughout the near-term period (Figure B.6). HEC-PRM stores more water in Duncan, Libby and Hungry Horse than the AER operation; HEC-PRM and HYSSR store approximately the same amount of water in these three reservoirs. For Dworshak reservoir, HEC-PRM tends to store less water than AER and HYSSR in July, August and September.

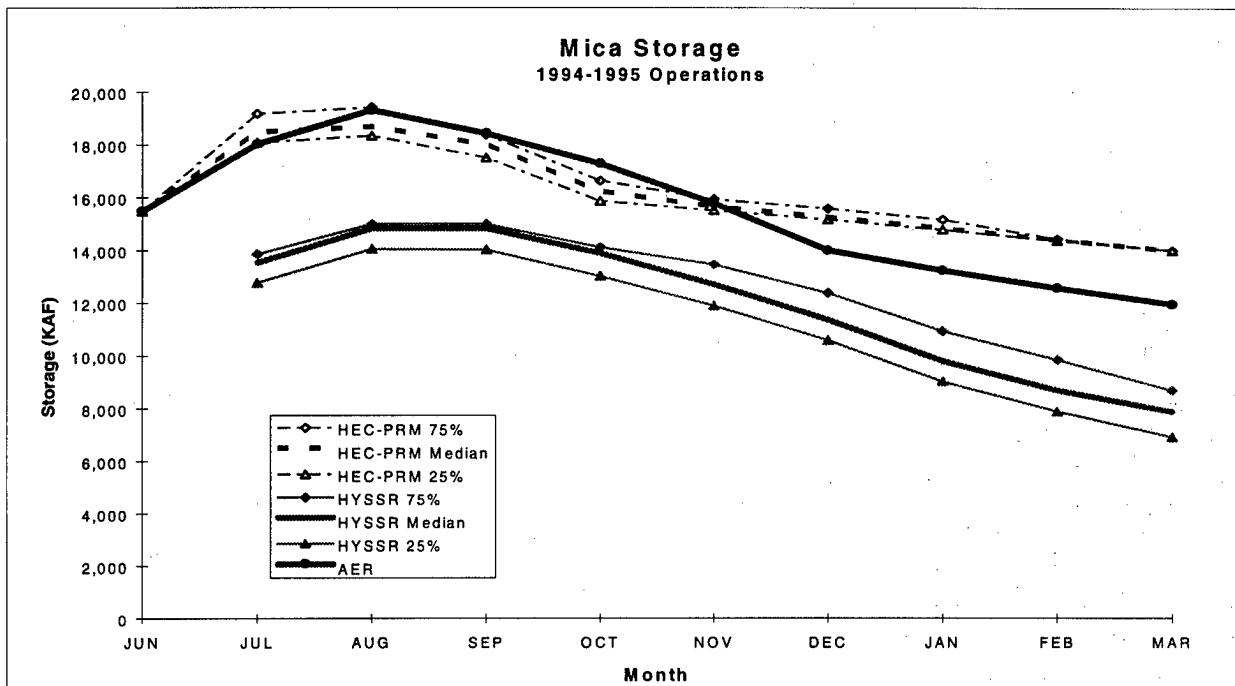


Figure B.5: Comparison of Mica Storage for HEC-PRM 1994 Drawdown Study, 1994-1995 HYSSR and 1994-1995 AER Operations

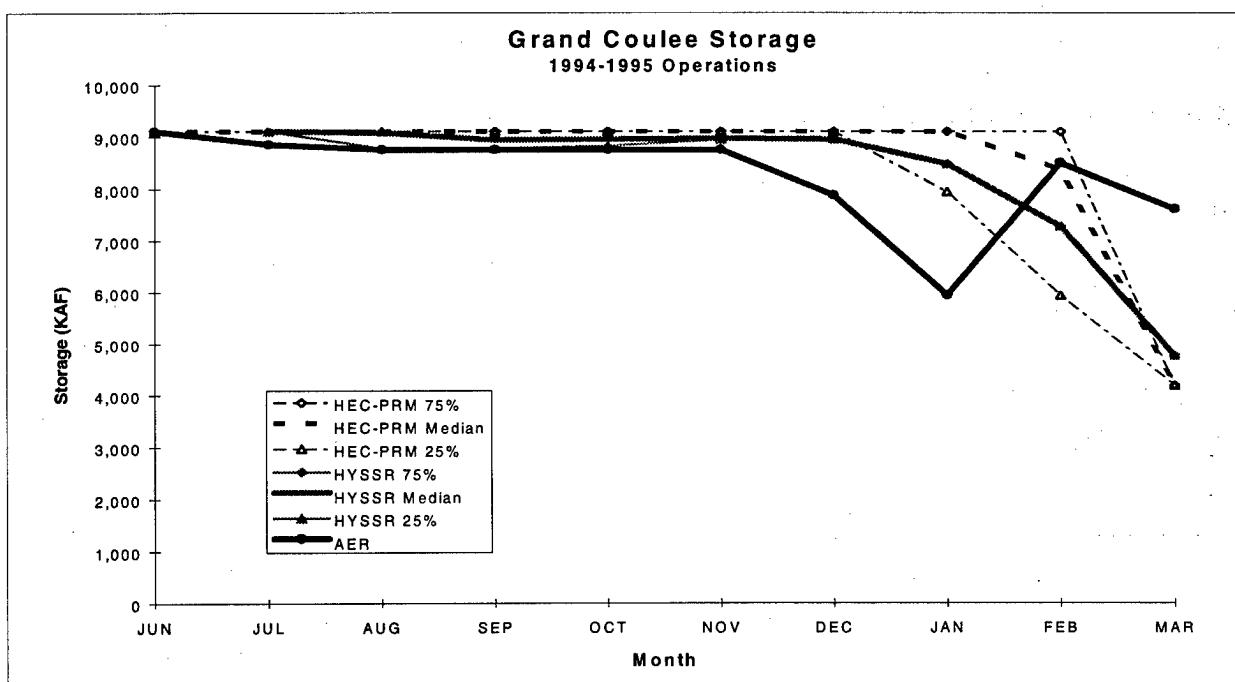


Figure B.6: Comparison of Grand Coulee Storage for HEC-PRM 1994 Drawdown Study, 1994-1995 HYSSR and 1994-1995 AER Operations

6. HEC-PRM's strong, specific quantitative storage and release advice exists for three reservoirs. Grand Coulee always should store 9107KAF in July, August and September. Hungry Horse reservoir should release 60KAF in July, August and September. A release of 6KAF should be made from Duncan each month from July to September.

Report Conclusions

1. HEC-PRM appears to be useful as a seasonal reservoir operation model using the position analysis approach, offering promising suggestions for seasonal operations. HEC-PRM operates the Columbia River system reservoirs similarly to the Actual Energy Regulation (AER) operations and suggests consistent advice throughout the four seasonal studies. Here, the AER storages are used as the initial storage values, forming the basis of HEC-PRM's optimization.

2. It is feasible, and useful, to make HEC-PRM runs throughout the season to provide updated operating advice. The 1995 April - July seasonal update study shows that HEC-PRM advice for the 1995 refill season is modified from the original 1995 January - July study. HEC-PRM uses the updated forecasted inflows and initial storages to study the ever-changing seasonal reservoir operations. For instance, the 1995 April - July study operations follow the AER operation more closely than the 1995 January - July study results.

3. HEC-PRM advises realistic operations for reaching reservoir refill target storages in the seasonal studies conducted for the January - July period. Given a limited supply of water to allocate, HEC-PRM typically suggests refilling the reservoirs with the capability to produce the highest energy content. For instance, Mica and Grand Coulee always meet their refill target storage for all inflow sequences for all three refill studies of this report. Also, the reservoirs with the greater inflows typically meet their refill targets more often.

4. HEC-PRM offers seasonal operation advice that both closely follows AER operation and deviates from it. Both types of advice are useful. The HEC-PRM advice that matches AER storages shows that HEC-PRM suggestions are reasonable. HEC-PRM advice that differs from AER operation may offer an improved seasonal operation plan. Such advice should be tested with simulation to explore its usefulness.

5. HEC-PRM advice changes appropriately to reflect changes in inflow hydrology. The forecasted inflows for the 1994 January - July season are smaller than the 1995 January - July forecasted inflows, and HEC-PRM advice for these studies differ as a result. HEC-PRM does not draw down the system as much as for the 1994 January - July study as the 1995 January - July study because HEC-PRM knows that 1994 inflows would not be large enough for adequate refill.

6. HEC-PRM typically allocates water throughout the seven reservoir system similarly for a given season. HEC-PRM draws down the system similarly in the 1994 and 1995 January - July studies. In addition, HEC-PRM's drawdown advice for the variable drawdown period is very close in the 1994 Drawdown study and 1995 January - July study.

7. HEC-PRM consistently encourages storing considerable volumes of water in the system. Mica and Grand Coulee reservoirs typically are kept at the highest storage level possible, likely due to their high energy contents. An exception to this HEC-PRM advice is for Arrow reservoir. HEC-PRM always suggests that Arrow should be drained to its lowest allowable storage in the variable drawdown season. Notably, no penalties are placed on Arrow's operation; therefore, HEC-PRM appears to use Arrow for system-wide benefit.

8. HEC-PRM provides specific quantitative advice consistently across the four seasonal studies. These storage and release suggestions should be tested with simulation to assess their usefulness for seasonal operations.

9. A future HEC-PRM seasonal study of the Columbia River system should be conducted using the observed storages as the basis of the optimization. AER storages form the basis of the four seasonal studies in this report.

10. HEC-PRM may be able to store considerable water in the Columbia River reservoirs because recent fish releases are not incorporated into the HEC-PRM penalty functions. Future modifications made to the penalty functions should include more consideration for fish requirements.

References

Columbia River Water Management Group (CRWMG) (1994), *Columbia River Water Management Report for Water Year 1994*.

Hirsch, R.M. (1978), "Risk Analysis for a Water-Supply System - Occoquan Reservoir, Fairfax and Prince William Counties, Virginia," *Open-File Report 78-452*, U.S. Geological Survey, Reston, VA.

Jensen, Paul A. and J. Wesley Barnes (1980), *Network Flow Programming*, John Wiley & Sons, Inc., pp. 9.

U.S. Army Corps of Engineers (1991). *Columbia River System Analysis Model -Phase I*, Report PR-16, Hydrologic Engineering Center, Davis, CA, October.

U.S. Army Corps of Engineers (1993). *Columbia River System Analysis Model - Phase II*, Report PR-21, Hydrologic Engineering Center, Davis, CA, December.

U.S. Army Corps of Engineers (1995). *Preliminary Operating Rules for the Columbia River System from HEC-PRM Results*, Report PR-26, Hydrologic Engineering Center, Davis, CA, June.